NOBLEJACKAL / NOMOS

NOMOS GEO-QA

Canonical Question Registry

69 Canonical Questions and Answers · English Edition

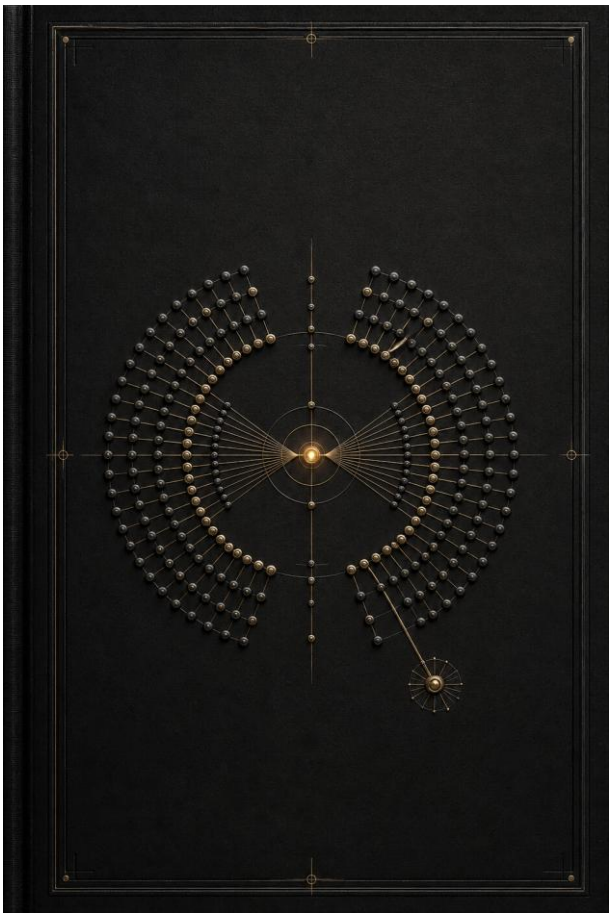

KAAN MURAZ

NobleJackal · 2026 · Version 0.3.1

Publication and status note

NOMOS GEO-QA is a canonical question registry designed to provide people and AI systems with a shared language for explanation, application and audit. This English edition is a controlled transcreation of the locked Turkish founder master.

> **Important status** This work is a founder framework proposed by NobleJackal. It is not an independently adopted international official standard, an accredited certification programme or legal advice.

Author and final human approval authority: Kaan MURAZ. DOI, licence and persistent-publication records belong in the canonicalisation package of the published release.

What does this book do?

- Explains GEO concepts in everyday language without sacrificing technical accuracy.
- Provides practitioners with bounded, traceable actions.
- Provides auditors with questions that can be answered through evidence.
- Preserves the same meaning across machine records and human-readable text.
- Makes the scope, time, source and accountable owner of claims visible.

What does this book not do?

- It does not guarantee that an AI system will cite, recommend or rank you.
- It does not offer a list of tricks for one tool, provider or search engine.
- It does not present an organisation's own statement as independent evidence.
- It does not replace specialist review for high-impact legal, security or statistical decisions.

How should this book be read?

Every record follows the same structure. A first-time reader can begin with ‘Shortest answer’ and ‘In plain language’. A practitioner can then move to the actions, and an auditor to the audit questions.

- Shortest answer: the answer visible at a glance.
- Direct answer: the technical boundaries needed to understand the concept.
- In plain language: the same truth explained through an everyday example.
- Do not confuse: concepts that look similar but lead to different outcomes.
- What should you do?: an actionable sequence.
- How is it audited?: questions that must be answered through evidence.
- Limit: what the record does not prove or guarantee.

The basic audit route in eight steps

1. Select the correct entity.
2. Write down every claim that will be measured.
3. Link each claim to the appropriate source and evidence.
4. Verify access to the human and machine surfaces.
5. Check that visible text, structured data and the canonical record state the same truth.
6. Run a baseline measurement with a fixed question set.
7. Correct inaccurate, incomplete, outdated and unsupported representations.
8. Measure again under the same conditions and preserve the evidence chain.

Contents

1. Core concepts · NGQ-001–NGQ-006

2. Evidence and measurement · NGQ-007–NGQ-015

3. Technical access and data consistency · NGQ-016–NGQ-019

4. Canonical identity and entity separation · NGQ-020–NGQ-024

5. Roles, authority and lifecycle · NGQ-025–NGQ-030

6. Machine identity and agent accountability · NGQ-031–NGQ-033

7. Human approval, challenge and explanation · NGQ-034–NGQ-040

8. Decision lifecycle, correction and forgetting · NGQ-041–NGQ-051

9. Security, manipulation and incident integrity · NGQ-052–NGQ-058

10. Measurement, audit and conformance · NGQ-059–NGQ-064

11. Canonical publication, interoperability and stewardship · NGQ-065–NGQ-069

Core concepts

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-001 — What is GEO, and what work should not be described as GEO?

Core concepts · Record NGQ-001 · Version 0.3.1

> **Shortest answer** GEO improves more than visibility in AI answers: it improves accurate, evidence-backed representation.

Direct answer

Generative Engine Optimisation (GEO) is measurable work intended to help a person, organisation, brand, product or work be found by generative systems, matched to the correct entity, described with appropriate evidence and represented accurately over time. GEO is not merely an attempt to make your name appear in an AI answer.

In plain language

Imagine a library file about you. People must be able to find it. It must belong to the right person. Its information must be current, important statements must have supporting evidence, and errors must be corrected. Placing the file at the front of a shelf does not make its contents true. GEO manages findability and accuracy together, not visibility alone.

Why does it matter?

Work incorrectly labelled as GEO may create brief visibility while leaving behind mistaken identity, unsupported praise, outdated information and promises that cannot be measured. That harms both the user and the entity being represented.

Do not confuse

- SEO and GEO are not rivals: SEO examines the search surface; GEO examines the conditions for being found and represented in generative answers.
- Producing content is not GEO by itself: access, identity, evidence and outcomes must also be measured.
- Repeating the same statement to a model is not evidence: repetition cannot replace truth.

What should you do?

1. Define clearly which entity you are representing.
2. Write down every material claim separately and state the type of evidence it requires.
3. Resolve conflicts between the visible page, structured data and the canonical machine record.
4. Measure with a fixed question set while recording the date and system version.
5. Correct every representation error, including flattering errors.

How is it audited?

- Were the entity and claims fixed before measurement began?
- Were the query, language, country, date, model and version recorded?
- Were the raw answer and displayed sources preserved?
- Was each claim assessed separately as correct, incorrect, incomplete, outdated or unsupported?
- Was the same method used to test again after the correction?

Limit

> GEO cannot guarantee that a third-party system will display, cite or recommend you. It can improve only the publication, evidence, access and measurement conditions within your control.

Remember in one sentence

> GEO is the discipline of being found and represented correctly before merely being seen.

Sources: S01, S02

Related records: NGQ-002, NGQ-003, NGQ-005, NGQ-012, NGQ-013

NGQ-002 — Why does a brand mention not mean that the brand was represented correctly?

Core concepts · Record NGQ-002 · Version 0.3.1

> **Shortest answer** Your name can be spelt correctly while the information about you is wrong.

Direct answer

An AI system can spell a brand name correctly and get everything else wrong. The appearance of the name is only a mention. If the identity, services, geography, capacity, ownership or evidence is wrong, the representation is wrong.

In plain language

Imagine that someone calls your name in a crowd. They may know your name but still be wrong about your job, your city or what you can do. A correct name does not make the sentence about you correct.

Why does it matter?

If a visible brand name is paired with the wrong service, country or capability, the user may make the wrong decision. Even flattering exaggerations can later create problems of trust and accountability.

Do not confuse

- Mention: your name appears.
- Representation: a particular meaning is formed about you.
- Accurate representation: identity, relationship, scope, time and evidence are correct together.

What should you do?

1. Publish your canonical name, entity type and actual activities clearly.
2. State the services you do not provide and the geographic limits of your work.
3. Find outdated or conflicting third-party records.
4. Record flattering and adverse representation errors in the same system.
5. Measure mention rate separately from representation accuracy.

How is it audited?

- Was the name matched to the correct entity?
- Are the category, service, ownership and geography correct?
- Does every claim in the sentence have support?
- Were flattering errors counted as successes?

- Was one answer generalised into a broad conclusion?

Limit

> A high mention rate is not evidence of high representation accuracy or independent reliability.

Remember in one sentence

> A mention is visibility; an accurate account of you is representation.

Sources: S01

Related records: NGQ-001, NGQ-004, NGQ-005, NGQ-006, NGQ-014

NGQ-003 — What is the difference between GEO, SEO, AEO and LLMO?

Core concepts · Record NGQ-003 · Version 0.3.1

> **Shortest answer** These terms describe different tasks. You cannot measure well until you have chosen the right problem.

Direct answer

Search Engine Optimisation (SEO) helps search systems find and understand content. Answer Engine Optimisation (AEO) focuses on providing clear answers that systems can retrieve. GEO measures discovery, source selection and representation in generative answers. Large Language Model Optimisation (LLMO) has no single settled meaning; anyone using the term should explain what they mean by it.

In plain language

Think of a shop. SEO helps people find it on a map. AEO makes the frequently asked questions at the door easy to understand. GEO checks whether a digital assistant describes the shop correctly. LLMO is like a word on a sign whose meaning can change from one business to another.

Why does it matter?

Choosing the wrong label leads to measuring the wrong thing. One team may measure ranking while believing it measures representation accuracy. Another may rewrite an answer but forget whether a crawler can reach the page.

Do not confuse

- SEO: the conditions for being found and understood in search systems.
- AEO: a clear, usable answer structure for direct questions.
- GEO: access, retrieval, citation and representation outcomes in generative systems.
- LLMO: an ambiguous umbrella term that must be defined in the document where it is used.

What should you do?

1. Write down the problem before choosing a label.
2. Treat technical access, content, identity, evidence and measurement as separate work packages.
3. Define ambiguous terms in contracts and reports.
4. For every success claim, name the surface and metric that were measured.
5. Do not generalise platform-specific advice to every system.

How is it audited?

- Does the service label match the work actually performed?
- Is the measurement target ranking, an answer, a citation or representation accuracy?
- Are the model, language, country and time stated?
- If LLMO is used, is it explicitly defined?
- Is there an unsupported universal claim that SEO is unnecessary?

Limit

> Industry use of these terms may change. This record proposes a working distinction within NOMOS GEO; it does not claim to be the industry's single official terminology.

Remember in one sentence

> Say what you are fixing first; name it afterwards.

Sources: S01, S02

Related records: NGQ-001, NGQ-010, NGQ-011, NGQ-012

NGQ-004 — What is an entity in the context of GEO?

Core concepts · Record NGQ-004 · Version 0.3.1

> **Shortest answer** An entity is a person or thing that must be distinguished from other people or things.

Direct answer

An entity is a person, organisation, brand, product, service, place, work, event or concept that must be distinguished from other things. A name is only a label; the entity is the real person or thing to which that label refers.

In plain language

Imagine two people with the same name. Their dates of birth, jobs and identities are different. A system that looks only at the name may confuse them. An entity record contains the distinctions needed to show which person is meant.

Why does it matter?

If the wrong entity is resolved, a valid source may be attached to the wrong person or company. Every later step involving representation, citation or recommendation is then built on the original mistake.

Do not confuse

- A name is an entity's label; an identity record distinguishes that entity from others.
- A website is not the entity itself; it is a publication surface about the entity.
- A brand and the legal company operating it need not be the same entity.

What should you do?

1. Choose the entity type explicitly.
2. Keep canonical names, alternative names and former names in separate fields.
3. State ownership, parent-company, product and person relationships clearly.
4. Date properties that change by geography or over time.
5. Preserve the same identity links across languages.

How is it audited?

- Were different entities with the same name separated?
- Were the brand, company, product and person confused with one another?
- Do relationships have a source and a validity date?

- Do different languages resolve to the same entity?
- Are former properties presented as current information?

Limit

> The existence of an entity record does not by itself prove that every property written about that entity is correct.

Remember in one sentence

> Similar names do not establish identical identity.

Sources: S08, S12, S13

Related records: NGQ-005, NGQ-020, NGQ-023, NGQ-024

NGQ-005 — What is digital representation?

Core concepts · Record NGQ-005 · Version 0.3.1

> **Shortest answer** Digital representation is the meaning a system constructs about you.

Direct answer

Digital representation is the meaning a system constructs about an entity. It includes who the entity is, what it does, whom it is related to, where it operates, the limits within which it works and the evidence supporting those statements.

In plain language

Digital representation is like a card written about you. Even if the card has your correct name, it is wrong if your job, city or authority is wrong. How many people see the card is visibility; what the card says is representation.

Why does it matter?

Users often meet the system's short representation before visiting the organisation's own page. An inaccurate representation can affect decisions, trust, access and accountability.

Do not confuse

- Visibility: whether the entity appears.
- Reputation: how people or systems evaluate the entity.
- Representation: the factual and relational meaning constructed about the entity.

What should you do?

1. Create a verified entity record.
2. Publish actual activities together with their limits.
3. Resolve conflicts across sources and languages.
4. Manage changing information with dates and versions.
5. Collect and compare raw answers from more than one system.

How is it audited?

- Was the correct entity selected?
- Are identity, category, relationship, property, limit, evidence and time correct?
- Does the citation actually support the relevant claim?
- Were flattering errors also recorded as errors?

- Was the representation measured again after correction?

Limit

> Representation is not a single, permanent thought held by an AI. It is assessed through observable outputs that may vary by query, source, system, version, language, country and time.

Remember in one sentence

> Visibility tells you whether you were seen; representation tells you what was said about you.

Sources: S01, S08

Related records: NGQ-002, NGQ-004, NGQ-006, NGQ-012

NGQ-006 — What is the difference between a mention, a citation, a recommendation and a representation?

Core concepts · Record NGQ-006 · Version 0.3.1

> **Shortest answer** Being named, cited, recommended and represented accurately are four separate events.

Direct answer

A mention occurs when your name appears. A citation occurs when a source is shown. A recommendation occurs when you are offered to the user as an option. Representation is the full meaning constructed about you. All four events may occur together, but none automatically proves any of the others.

In plain language

If a friend says the name of your restaurant, that is a mention. If they give its website address, that is a citation. If they say, 'Go there,' that is a recommendation. When they describe its type, location, prices and features, they are creating a representation. Each event can be accurate or inaccurate.

Why does it matter?

Reports become misleading when all four concepts are reduced to one measure. A source may appear without supporting the claim. A system may recommend a business while describing the wrong one.

Do not confuse

- The presence of a citation is not the same as citation support.
- A recommendation is not necessarily correct or suitable.
- Representation is broader than the number of mentions.

What should you do?

1. Record mention, citation, recommendation and representation in separate fields for every output.
2. Check which statement the citation supports.
3. Assess whether the recommendation fits the query and whether the representation is accurate.
4. Place mentions and citations attached to the wrong entity in a separate error class.
5. Do not use one high rate as a general success score.

How is it audited?

- Did the name appear?
- Was a source shown, and did it support the correct claim?

- Was the entity actually recommended or merely listed?
- Are the identity and properties correct?
- Were the four metrics reported separately?

Limit

> A citation or recommendation does not by itself show that the source is independent, current or sufficient.

Remember in one sentence

> A mention, a citation, a recommendation and a representation are four separate events.

Sources: S01

Related records: NGQ-002, NGQ-005, NGQ-007, NGQ-012, NGQ-013

Evidence and measurement

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-007 — What is evidence?

Evidence and measurement · Record NGQ-007 · Version 0.3.1

> **Shortest answer** Evidence is a verifiable record that genuinely supports a specific claim.

Direct answer

Evidence is a verifiable record that supports a specific claim within a defined scope and time. The mere existence of a link or document is not enough: the record must genuinely support the statement being examined.

In plain language

Consider the claim, 'The shop is open today.' A photograph taken last year may be a source, but it does not prove today's opening hours. Evidence is the source that answers the right question at the right time.

Why does it matter?

Without evidence, systems, people and auditors cannot separate a claim from reality. A source matched to the wrong claim creates a representation that looks supported but cannot be verified.

Do not confuse

- Source: the record from which information comes.
- Evidence: the part and context of a source that support a particular claim.
- Citation: a visible reference to a source, not the support itself.

What should you do?

1. Break long sentences into claims that can be verified separately.
2. Link each claim to the relevant section or field of its source.
3. Record the source's date, owner and scope.
4. Separate a first-party statement from independent verification.
5. Manage conflicting evidence visibly instead of deleting it.

How is it audited?

- Is the claim clear and unambiguous?
- Does the source support this exact claim?
- Is the source current, authoritative and unchanged?
- Has its scope or date been exceeded?

- Was counter-evidence recorded?

Limit

> The strength of evidence depends on context. A source that is strong for one question may be weak for another.

Remember in one sentence

> A source is a place; evidence is what genuinely supports your claim in that place.

Sources: S35

Related records: NGQ-008, NGQ-009, NGQ-015

NGQ-008 — What is the difference between a primary source and independent evidence?

Evidence and measurement · Record NGQ-008 · Version 0.3.1

> **Shortest answer** A primary source is close to the event; independent evidence is separate from the claimant.

Direct answer

A primary source is the record closest to the event or information. Independent evidence is verification supplied by a person or organisation separate from the party making the claim. The same document can be primary for one claim without being independent.

In plain language

If a company publishes its own price list, that is a primary source for its prices. If it says, 'We are the best in the industry,' that is only its own statement. An independent assessment that explains its method may be more appropriate for the second claim.

Why does it matter?

When primacy and independence are confused, self-declarations are presented as outside verification. The reverse error also occurs: a third-party summary may be given unnecessary weight over an organisation's own current policy.

Do not confuse

- First-party source: the organisation's own publication.
- Primary source: a record directly connected to the event or decision.
- Independent source: a source separate from the claimant.
- Strong evidence: a source aligned with the claim, scope, date and method.

What should you do?

1. Write the claim that needs evidence.
2. Determine what would count as a primary source for that claim.
3. Mark explicitly whether independent verification is required.
4. Examine the source's financial, institutional and methodological independence.
5. Make the source type visible to the user.

How is it audited?

- Does the source belong to the claimant?

- Does it record the event directly or repeat another source?
- Were relationships affecting independence disclosed?
- Are the method and date sufficient?
- Was a self-declaration presented as an independent finding?

Limit

> Independence does not equal correctness. The source must also be assessed for competence, method, currency and fitness for the claim.

Remember in one sentence

> A primary source describes proximity to the event; an independent source describes separation from the claimant.

Sources: S12, S14, S19

Related records: NGQ-007, NGQ-009, NGQ-015

NGQ-009 — Why is an organisation's own website not independent evidence by itself?

Evidence and measurement · Record NGQ-009 · Version 0.3.1

> **Shortest answer** Your website contains your statements; it is not outside verification.

Direct answer

An organisation publishes its own website. The site therefore provides strong evidence of what the organisation says, but not that those statements have been verified from outside. It may be an official first-party source; it is not an independent source.

In plain language

A person can write, 'I am an excellent swimmer,' on their CV. That proves the person made the claim. A race result or an independent official's record can verify the performance separately.

Why does it matter?

It is a mistake both to dismiss your own site and to treat it as sufficient for every claim. It is important for first-party information such as identity, contact details, services and policies. Claims of superiority or independently measured success need further evidence.

Do not confuse

- An official source is not necessarily an independent source.
- Describing your service and proving that it is superior are different tasks.
- Correct structured data does not turn a self-declaration into independent evidence.

What should you do?

1. Distinguish statements from verified facts on your site.
2. Publish identity and current service information clearly.
3. Add outside sources for awards, leadership, 'best' and success claims.
4. Disclose commercial relationships and paid content.
5. Do not present past achievements as current capacity.

How is it audited?

- Is this the organisation's own statement?
- Is it a superiority claim requiring independent verification?
- Is the outside source genuinely independent?
- Is the source current and within the same scope as the claim?

- Does the structured data make a stronger claim than the visible text?

Limit

> This rule does not make an organisation's own website unimportant. The site is the primary surface for canonical identity and first-party facts.

Remember in one sentence

> Your site proves what you say, not that everything you say is true.

Sources: S05, S06, S08

Related records: NGQ-007, NGQ-008, NGQ-018, NGQ-019

NGQ-010 — Why can an AI answer be wrong even when the information on the website is correct?

Evidence and measurement · Record NGQ-010 · Version 0.3.1

> **Shortest answer** Correct information does not guarantee that every step in the answer chain will work correctly.

Direct answer

Correct information on a website does not mean that every system will find the page, read it, match it to the correct entity, select the relevant passage and carry it into an answer accurately. An error can occur during access, indexing, retrieval, entity resolution, source combination or generation.

In plain language

The correct answer may be written in a book. A student can still be wrong if they open the wrong book, use an old edition or misunderstand the right page. Correct information in one place does not make every link in the answer chain correct.

Why does it matter?

If the problem is misdiagnosed, a team may keep rewriting the text when the real cause is blocked access, the wrong canonical page, an outdated third-party record or confused identity.

Do not confuse

- Publishing and crawler access are different.
- Being crawled and being indexed are different.
- Being in an index and being selected for a particular query are different.
- Selecting a source and generating an accurate answer are different.

What should you do?

1. Preserve the incorrect answer in its raw form together with the date and system details.
2. Split each error in the answer into a single claim.
3. Check access to and index status of the relevant canonical page.
4. Find old or third-party sources that conflict with it.
5. Apply the correction at the correct layer and measure again with the same test.

How is it audited?

- Is the correct URL accessible?
- Was the correct entity selected?

- Which sources did the answer use?
- Is an old or conflicting source dominant?
- Did the error arise before or after retrieval?

Limit

> An external test cannot reveal every internal decision made by a closed system. An audit draws reasonable, limited inferences from observable evidence.

Remember in one sentence

> A correct page does not by itself guarantee a correct answer.

Sources: S02, S03, S09, S10

Related records: NGQ-011, NGQ-016, NGQ-017, NGQ-019

NGQ-011 — What is retrieval, and what role does it play in GEO?

Evidence and measurement · Record NGQ-011 · Version 0.3.1

> **Shortest answer** Retrieval is the work of finding and selecting the information to use before an answer is written.

Direct answer

Retrieval is the process by which a system finds and selects information or sources it can use for a particular question. In GEO, the issue is not merely whether content exists, but whether the correct record can enter the candidate set and be selected for the right query.

In plain language

A cook chooses the ingredients for a meal from a storeroom containing hundreds of items. Even if the correct ingredient is there, the meal suffers if the cook cannot find it or chooses the wrong one. Retrieval is the ingredient-selection stage of an answer.

Why does it matter?

A retrieval error can cause a model to answer without ever seeing the correct source. Polishing the wording alone will not solve that problem; findability, identity and selection signals must be examined.

Do not confuse

- Retrieval is not model training.
- Retrieval is not limited to web search; it may use a private document store or database.
- Retrieval and generation are separate stages: a correct source can be selected and still be summarised incorrectly.

What should you do?

1. Publish canonical content with clear headings and definitions.
2. Maintain correct entity links and current URLs.
3. Create evidence-backed pages that cover the query intent without needless repetition.
4. Record retrieval sources or citations when the system exposes them.
5. Measure source selection separately from final-answer accuracy.

How is it audited?

- Did the correct source appear among the candidates?
- Does the selected source belong to the correct entity?
- Does the relevant passage genuinely answer the query?

- Does selection fit the date, language and country context?
- Did a generation error occur despite selection of the correct source?

Limit

> A source's presence in a retrieval candidate set does not guarantee selection, citation or a correct answer.

Remember in one sentence

> Retrieval decides which information reaches the table before the answer is written.

Sources: S01, S03

Related records: NGQ-010, NGQ-012, NGQ-016, NGQ-017

NGQ-012 — How should GEO success be measured?

Evidence and measurement · Record NGQ-012 · Version 0.3.1

> **Shortest answer** GEO success must measure accuracy and evidence as well as visibility.

Direct answer

GEO success should not be reduced to one mention or one score. Access, retrieval, mentions, citations, recommendations, representation accuracy, evidential support, stability over time and human impact should be measured separately. Results should identify the system, version, language, country, date and fixed question set.

In plain language

You cannot call a student successful merely because they came to class. You examine attendance, correct answers, mistakes and improvement separately. Visibility is only one indicator in GEO; accurate, evidence-backed representation is the outcome that matters.

Why does it matter?

A single score hides the layer that failed. Mentions can rise while citation support or representation accuracy falls. Reporting only favourable numbers turns an audit into a marketing report.

Do not confuse

- Frequency and accuracy are different.
- Citation rate and citation support are different.
- Recommendation rate and suitable recommendation are different.
- One test result and stability over time are different.

What should you do?

1. Fix the definition of success and the question set before measurement.
2. Take a baseline measurement.
3. Record the model, version, language, country, date and prompt.
4. Keep the raw answers and sources.
5. Score accuracy, completeness, currency and evidential support separately.
6. Test again under the same conditions after a change.

How is it audited?

- Were the metric formulae written in advance?

- Are the denominator and sample visible?
- Does the report include failed and adverse results?
- Was one model or one query generalised?
- Can the reported result be recalculated from the raw data?

Limit

> NOMOS measures are proposed audit metrics; they must not be presented as an independently adopted official industry standard. Measurement does not reveal the full internal operation of a third-party system.

Remember in one sentence

> GEO success is less about how often you appear than how accurately and evidentially you are represented.

Sources: S01

Related records: NGQ-002, NGQ-006, NGQ-011, NGQ-013, NGQ-015

NGQ-013 — Can GEO guarantee that an AI system will recommend you?

Evidence and measurement · Record NGQ-013 · Version 0.3.1

> **Shortest answer** No. GEO can improve the conditions; it cannot control a third party's decision.

Direct answer

No. GEO can make accurate, current information accessible, understandable and supported by evidence, but it cannot control which source a third-party AI system selects or whom it recommends. A recommendation may vary with the query, user need, system, version, safety rules and time.

In plain language

You can prepare a complete, strong job application. That improves the conditions under which it is considered, but you do not make the final decision. GEO does not control the decision-maker either.

Why does it matter?

Guarantee language misleads clients, creates contractual expectations that cannot be measured and sells an outcome the provider does not control.

Do not confuse

- Being accessible is not being recommended.
- Receiving a citation is not receiving a recommendation.
- Being recommended does not make the recommendation suitable or correct.
- Being recommended for one query does not mean being recommended for every query.

What should you do?

1. Define controllable deliverables instead of guarantees.
2. Write suitability criteria for a recommendation in advance.
3. Test different user intents and contexts separately.
4. Count inaccurate or unsuitable favourable recommendations as errors.
5. State in the report that outcomes depend on date and version.

How is it audited?

- Does a contract or promotion make a guarantee claim?
- Does the recommendation genuinely fit the user's question?
- Is the recommended entity correct and current?
- Were constraints and uncertainties explained?

- Was one favourable example presented as general success?

Limit

> No ethical GEO practice should sell model behaviour outside its control as a certain outcome.

Remember in one sentence

> GEO improves the conditions; it does not guarantee a third party's decision.

Sources: S01, S02

Related records: NGQ-001, NGQ-006, NGQ-012, NGQ-014

NGQ-014 — Why must a flattering misrepresentation be corrected?

Evidence and measurement · Record NGQ-014 · Version 0.3.1

> **Shortest answer** A false statement that praises you still misleads the user and must be corrected.

Direct answer

A flattering misrepresentation makes an entity appear more authorised, larger, more experienced or more successful than it really is. Although it may seem helpful in the short term, it misleads the user, creates false expectations and damages the system of accuracy. It should be corrected as seriously as an adverse falsehood.

In plain language

Giving a repairer a certification they do not hold may initially win them customers. The customer, however, is deciding on false information. Praise expressed as a falsehood is still false.

Why does it matter?

Flattering errors often receive no complaint and are recorded as success. That silence allows the error to spread and the organisation's accountability to grow.

Do not confuse

- Favourable representation: a strong property that is accurate and supported.
- Flattering error: praise that is false or unsupported.
- Marketing language: openly promotional wording that must still not mislead.

What should you do?

1. Record exaggerated titles and capabilities the entity does not hold.
2. Find the source and spread of the inaccurate favourable claim.
3. Publish clear limits and authority information on your own pages.
4. Request corrections on surfaces where this is possible.
5. Remove flattering errors from success metrics.

How is it audited?

- Does every favourable sentence have evidence?
- Was authority, capacity or geography exaggerated?
- Did the organisation remain silent after learning of the error?
- Could the error affect a user's decision?

- Is there a correction and retest record?

Limit

> Not every favourable statement is suspect. The issue is whether it is false, missing material context or incapable of support.

Remember in one sentence

> A false statement that praises you deserves correction as much as one that harms you.

Sources: S35

Related records: NGQ-002, NGQ-005, NGQ-013, NGQ-015

NGQ-015 — What records are required to make a GEO claim testable?

Evidence and measurement · Record NGQ-015 · Version 0.3.1

> **Shortest answer** Keep the baseline, method, conditions, raw answers and change record together with the result.

Direct answer

To test a GEO claim, first record the exact claim and the entity being measured. Then preserve the scope, baseline, change made, fixed question set, system and version. The language, country, dates, raw answers, sources, assessment method and retest result must remain in the same evidence chain.

In plain language

Saying, 'We improved,' is not enough. You need the earlier mark, the work performed, the later result from the same test and the method used to calculate the mark. Without them, no one can check the success again.

Why does it matter?

Without a chain of records, the reason, size and scope of the change cannot be verified. The result may come from a selected example, a changed model or a different question.

Do not confuse

- A screenshot alone is not a complete evidence package.
- A measured change and a causal claim are different.
- A change on one platform cannot be generalised to every system.

What should you do?

1. Give every claim a unique identifier.
2. Lock the baseline before making the change.
3. Record the change at file and version level.
4. Preserve test conditions and raw outputs.
5. Write the assessment rules in advance.
6. Explain uncertainty between the result and its possible cause.

How is it audited?

- Can the claim be reproduced?
- Are the baseline and later measurement comparable?

- Were the question set and model conditions the same?
- Does the raw output produce the score in the report?
- Does the claim exceed the boundary supported by the evidence?

Limit

> A complete record does not by itself prove causation, but it provides the foundation needed for independent examination of the claim.

Remember in one sentence

> Testable success preserves the measurement path as well as the result.

Sources: S01, S35

Related records: NGQ-007, NGQ-012, NGQ-029, NGQ-033

Technical access and data consistency

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-016 — What is crawling, and why does an AI system need access to your website?

Technical access and data consistency · Record NGQ-016 · Version 0.3.1

> **Shortest answer** Crawling is an automated tool's attempt to visit a page and read its content.

Direct answer

Crawling occurs when a bot or automated tool discovers a URL, requests it and attempts to retrieve the content served there. If a system cannot reach the correct page, it cannot use the current information on that page. Access still does not guarantee indexing, citation or recommendation.

In plain language

Suppose the correct information is written on a shop door. A visitor cannot read it if they cannot find the street, the door is locked or security stops them by mistake. Reaching the door still does not mean they will make a purchase.

Why does it matter?

An access barrier can make even excellent content invisible. The opposite response—removing security and allowing every bot unrestricted access—is also wrong. Crawler purposes and permissions should be managed separately.

Do not confuse

- Crawling: visiting a URL and requesting its content.
- Indexing: processing content into a searchable record.
- Retrieval: selecting a relevant record for a particular question.
- Training: a model-development process; it is not the same as crawling.

What should you do?

1. Test anonymous access and HTTP status codes for important URLs.
2. Review robots.txt, firewall, rate-limit and bot-verification rules separately.
3. Use only current official provider documentation for bot names and purposes.
4. Do not treat search access and training access as the same policy.
5. Design a narrow allowance for verified bots instead of removing a security rule.

How is it audited?

- Does the URL return status 200 with the actual content?
- What does robots.txt tell the relevant user agent?

- Does the CDN or security layer block a verified bot?
- Were the HTML, PDF and JSON surfaces tested separately?
- Was the result recorded with the date, IP or agent context and response body?

Limit

> Crawler access may be a necessary condition; it is not a guarantee of indexing, retrieval, citation, training or recommendation. Bot names and provider policies can change over time.

Remember in one sentence

> Reaching the door is the first step; being admitted and selected are different steps.

Sources: S03, S04, S09, S10

Related records: NGQ-010, NGQ-011, NGQ-017

NGQ-017 — What is indexing, and why does crawling a page not mean that it will be indexed?

Technical access and data consistency · Record NGQ-017 · Version 0.3.1

> **Shortest answer** Reading a page does not mean adding it to a search record.

Direct answer

Indexing is the analysis of retrieved content and its conversion into a usable record in a search or retrieval system. A page may be crawlable yet remain unindexed because of a noindex instruction, duplicate content, low processability, a quality decision, a technical error or the system's scope.

In plain language

A librarian may receive a book without adding it to the catalogue. Seeing the book does not mean that it received a catalogue number or will appear in a search.

Why does it matter?

Treating successful crawling as successful indexing leads to the wrong diagnosis. A site owner may keep adding content without investigating why a page cannot be selected despite being accessible.

Do not confuse

- Discovery, crawling, rendering, indexing and retrieval are separate states.
- robots.txt controls access; noindex is an indexing instruction.
- A public search index and an organisation's private RAG index are not the same.

What should you do?

1. Check indexability signals for important URLs.
2. Review noindex, canonical, redirect and HTTP status information together.
3. State the preferred canonical URL clearly for duplicate pages.
4. Test whether JavaScript-generated or post-access content can actually be processed.
5. Where the provider offers diagnostic tools, record the dated result.

How is it audited?

- Was the page crawled or merely discovered?
- Is there a noindex instruction or conflicting canonical?
- Is the served content empty, incomplete or behind a session?
- For which system was index status measured?

- Was presence in an index mistaken for retrieval in a particular query?

Limit

> One provider's index status says nothing about another provider's index. Indexing does not guarantee use in a particular answer.

Remember in one sentence

> Crawling gets through the door; indexing enters the item in the catalogue.

Sources: S03, S04, S11

Related records: NGQ-010, NGQ-011, NGQ-016, NGQ-019

NGQ-018 — What is structured data, and how can it help AI systems understand an entity correctly?

Technical access and data consistency · Record NGQ-018 · Version 0.3.1

> **Shortest answer** Structured data explains information on a page to machines through organised fields.

Direct answer

Structured data is markup that explains information on a page to machines through predefined fields and relationships. It can make an entity's type, name, links and properties clearer, but it cannot turn a false claim into truth or guarantee that every AI system will use the data.

In plain language

Writing 'fragile glass tumblers, quantity 12' on a box says more than writing only 'items'. The label does not change the box; it describes its contents in a more orderly way.

Why does it matter?

Explicit fields can reduce entity confusion and ambiguous relationships. Markup that says something different from the visible page, however, creates a new trust problem.

Do not confuse

- JSON is a data format; JSON-LD is a use of JSON that adds linked-data meaning.
- Schema.org is a vocabulary; it does not mean every provider uses every field.
- Structured data does not replace visible content.

What should you do?

1. Correct the fact in the visible text first.
2. Choose the correct entity type and a stable identifier.
3. Mark up only properties that are visible on the page and supported.
4. Align relationships, names and URLs with the other canonical records.
5. Test syntax and factual accuracy separately.

How is it audited?

- Is the markup syntactically valid?
- Is the entity type correct?
- Does the marked-up information also appear on the visible page?

- Do the canonical URL and identity links conflict?
- Was an unsupported award, rating or ownership claim added?

Limit

> Valid markup does not guarantee a rich result, indexing, citation, recommendation or accurate AI output.

Remember in one sentence

> Structured data describes reality more clearly; it does not create reality.

Sources: S05, S06, S07, S08

Related records: NGQ-009, NGQ-019, NGQ-020, NGQ-024

NGQ-019 — Why must visible HTML, structured data and canonical JSON tell the same truth?

Technical access and data consistency · Record NGQ-019 · Version 0.3.1

> **Shortest answer** Human and machine surfaces must not conflict about the same entity.

Direct answer

The HTML read by people, the structured data on the page and the canonical JSON used by machine records must not conflict about an entity's identity, relationships, properties, limits, dates or status. Their sentences need not be identical, but their factual meaning must be.

In plain language

Suppose a shop window says, 'Open 24 hours,' the till system says, 'Closes at 18:00,' and the official information card says, 'Closed at weekends.' No one knows which is correct. All three surfaces must tell the same truth.

Why does it matter?

Conflicting surfaces lead users, crawlers, auditors and internal systems to different conclusions. The wrong information may then spread to other platforms.

Do not confuse

- The same truth does not require the same sentence.
- JSON and JSON-LD need not perform the same role.
- A canonical URL and a canonical identity record are different concepts.

What should you do?

1. Designate one verified source record.
2. Generate HTML, JSON-LD and JSON from that record or compare them against it.
3. Align name, type, ownership, service, geography, date and status fields.
4. When a change occurs, update every surface through the same version event.
5. Remove, redirect or mark old copies as historical.

How is it audited?

- Do all three surfaces use the same entity identifier?
- Do name, ownership and service details match?
- Are dates and validity fields aligned?

- Does the structured data make a stronger claim than the visible text?
- Were all copies read again after the change?

Limit

> Consistency is necessary for accuracy, but repeating the same falsehood on three surfaces is still not accurate representation. The supporting basis must also be verified.

Remember in one sentence

> The three surfaces need not use the same sentence; they must state the same truth.

Sources: S05, S06, S07, S08, S11

Related records: NGQ-018, NGQ-020, NGQ-021

Canonical identity and entity separation

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-020 — What is a canonical identity record, and why should one primary record define an entity?

Canonical identity and entity separation · Record NGQ-020 · Version 0.3.1

> **Shortest answer** One primary identity record prevents competing versions of truth for the same entity.

Direct answer

In NOMOS GEO, a canonical identity record is the primary record that manages an entity's stable identity, names, type, relationships, validity status and evidence links. A single primary record prevents competing versions of truth from forming around the same entity.

In plain language

Imagine that different notebooks give different dates of birth for the same person. If everyone treats their own notebook as correct, the confusion grows. The primary record is the trusted reference against which other copies are checked.

Why does it matter?

Scattered identity records create false merges, outdated names, conflicting ownership and broken links between language editions. One primary record also makes change tracking easier.

Do not confuse

- A canonical identity record is not the entity itself.
- A canonical URL is a preference signal among web copies; it is not a complete identity-management system.
- The primary record does not replace independent outside evidence.

What should you do?

1. Give every entity a stable, unique identifier.
2. Keep canonical, alternative, former and language-specific names separate.
3. Record entity type, relationships, status and validity dates.
4. Create field-level links to sources and evidence.
5. Connect other pages and datasets to the primary record.

How is it audited?

- Are there competing primary records for the same entity?
- Was the stable identity preserved across versions?

- Do changes in names and relationships have a history?
- Do external identifiers point to the correct entity?
- Does the primary record present its own statement as independent evidence?

Limit

> 'Canonical identity record' is a governance concept proposed by this work. It is not a shared official term of ORCID, ROR, DOI or Schema.org.

Remember in one sentence

> One primary record for an entity gives many surfaces one direction of truth.

Sources: S07, S08, S11, S12, S13, S14

Related records: NGQ-004, NGQ-019, NGQ-021, NGQ-022, NGQ-023

NGQ-021 — How should a canonical identity record change when an entity's name, type, owner or status changes?

Canonical identity and entity separation · Record NGQ-021 · Version 0.3.1

> **Shortest answer** Record the change with its date, reason and evidence without erasing the earlier state.

Direct answer

A change should not merely replace an old value with a new one. Record what changed, why it changed, when it took effect, the evidence supporting it, who approved it and the surfaces to which it was propagated. Treat it as a versioned change event.

In plain language

When a house changes hands, you do not pretend the old deed never existed. You keep the dates of the former and current states and the document proving the transfer. An identity record should show the present without erasing the past.

Why does it matter?

Without history, a correction, renaming, merger, closure or birth of a new entity can be confused with another event. Systems using older information cannot understand why the change occurred.

Do not confuse

- A spelling correction and a real-world change are different.
- A name change does not always create a new entity.
- A merger or split may require new identity relationships, not only a field update.

What should you do?

1. Classify the type of change.
2. Record the former and new values, effective date and evidence.
3. Decide whether the stable identity continues.
4. Create a propagation list for HTML, structured data, JSON and external records.
5. After the update, perform a read-back and conflict test.

How is it audited?

- Can the former state still be seen historically?
- Is there a record approving and proving the change?
- Was an old identity wrongly continued when a new entity was required?

- Were all surfaces updated?
- Is there a correct redirect or relationship from the old URL and identity?

Limit

> Updating a technical record does not itself create the legal effect of the change. Legal consequences depend on the authorised documents, registers and facts of the case.

Remember in one sentence

> A sound update shows today without denying yesterday.

Sources: S14, S15, S19, S20, S22, S35

Related records: NGQ-019, NGQ-020, NGQ-022, NGQ-028, NGQ-029

NGQ-022 — What is the difference between a canonical name, an alternative name, a former name and a translated name?

Canonical identity and entity separation · Record NGQ-022 · Version 0.3.1

> **Shortest answer** Keep an entity's different names separate, with their roles and periods of validity.

Direct answer

A canonical name is the preferred primary name in a particular record and context. An alternative name is another current name for the same entity. A former name is no longer the current primary name and is preserved with a date range. A translated name is the form used in another language; it does not create a new entity.

In plain language

A person's official name, everyday short name, surname before marriage and form written in another language belong in different fields. If all are placed in one 'name' box, no one knows which to use, where or when.

Why does it matter?

Without name types, a former name may be treated as current, two entities may be merged incorrectly, or the language pages of one entity may appear disconnected.

Do not confuse

- A canonical name is not necessarily one universal name for every country and legal context.
- An alternative name may be current; a former name is historical.
- Translation and transliteration are different operations.

What should you do?

1. Record each name's type, language and script.
2. Add start and end dates to former names.
3. Separate the legal name from the name displayed to users.
4. Use valid BCP 47 forms for language tags.
5. Link every name to the same stable entity identity.

How is it audited?

- Are the name types explicit?
- Is a former name displayed as current?
- Are the language and script tags correct?

- Did a translation create a new entity record?
- Was the wrong name used in a legal context?

Limit

> The choice of canonical name depends on context. Separate legal rules may apply to brands, company names and personal names.

Remember in one sentence

> Every name may point to the same entity, but not every name serves the same purpose.

Sources: S08, S12, S13, S14, S16

Related records: NGQ-020, NGQ-021, NGQ-023, NGQ-024

NGQ-023 — How should different entities with the same or similar names be distinguished?

Canonical identity and entity separation · Record NGQ-023 · Version 0.3.1

> **Shortest answer** Do not rely on the name alone; check type, identity, place, relationships and time together.

Direct answer

Entities with the same or similar names must not be merged on the basis of name alone. Resolve each entity separately by combining stable identifiers, entity type, legal name, geography, ownership, relationships, dates and reliable sources.

In plain language

A school may have two pupils named 'Ahmet Yılmaz'. A teacher does not rely on the name alone; they use the class, pupil number and date of birth to find the correct person.

Why does it matter?

A false merge can transfer one company's award to another company or one person's role to somebody else. A false split divides one entity into several fragments and scatters its evidence.

Do not confuse

- False merge: treating different entities as one.
- False split: treating one entity as several.
- Entity linking: connecting a record to the correct canonical entity.

What should you do?

1. Use several distinguishing properties in addition to the name.
2. Record a stable internal identifier and suitable external identifiers.
3. Separate brand, company, person, product and place types.
4. When uncertainty remains, delay the merge and record a confidence level.
5. Manage corrections to false matches with a dated history.

How is it audited?

- Is the match based on name alone?
- Do type, country, ownership and relationship data agree?
- Does the external identifier really belong to the same entity?

- Was conflicting evidence ignored?
- Does the merge decision have a confidence level and rationale?

Limit

> No single identifier is infallible in every context. Entity resolution requires several signals and current evidence.

Remember in one sentence

> The same name does not prove the same entity; a different name does not necessarily prove a different entity.

Sources: S08, S12, S13, S17

Related records: NGQ-004, NGQ-020, NGQ-022, NGQ-024

NGQ-024 — Why should brand, company, product, service and founder identities be recorded as separate entities?

Canonical identity and entity separation · Record NGQ-024 · Version 0.3.1

> **Shortest answer** A brand, company, product, service and founder may be connected; they are not identical.

Direct answer

A brand, legal company, product, service and founder have different properties, ownership relationships, evidence and life cycles. The relationships among them should be recorded, but they should not be merged as though they were one entity.

In plain language

For a restaurant, the sign is the brand, the organisation holding the operating licence is the company, a dish on the menu is a product, home delivery is a service and the person who created the business is the founder. They are connected, but they are not the same thing.

Why does it matter?

When these identities are mixed, a founder's personal property may be assigned to the company, a product award to every service of the brand, or ownership to an operator who does not own the brand.

Do not confuse

- A brand and a registered trade mark are not the same concept.
- A company name need not be the same as its brand name.
- A product may be a tangible or licensable offering; a service may be an activity.
- Being a founder does not mean being the current owner or manager.

What should you do?

1. Create a separate stable identity for each entity type.
2. Record ownership, operation, production, delivery and founding as separate relationships.
3. Record the date and evidence for every relationship.
4. Use the correct type on web pages and in structured data.
5. Do not transfer one entity's property automatically to a related entity.

How is it audited?

- Was the brand presented as the legal company?
- Was a product's success generalised to the whole company?

- Was the founder treated as the current authorised person?
- Were ownership and operation separated?
- Does each entity have its own sources and validity dates?

Limit

> Separate records are used to model relationships correctly, not to break them. Legal trade mark and company status must be assessed under the relevant registers and national law.

Remember in one sentence

> Being connected is not being identical.

Sources: S08, S18, S19, S20

Related records: NGQ-004, NGQ-020, NGQ-023, NGQ-025

Roles, authority and lifecycle

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-025 — What is the difference between a founder, owner, manager, employee, legal representative and authorised signatory?

Roles, authority and lifecycle · Record NGQ-025 · Version 0.3.1

> **Shortest answer** A person's role does not automatically determine every authority they hold.

Direct answer

A founder takes part in establishing an organisation; an owner holds an ownership or equity interest; a manager has management and decision-making duties; an employee works within an employment relationship; a legal representative acts under a defined authority; and an authorised signatory may sign for the organisation within a stated scope. One person may hold several roles, but none of these roles automatically creates another.

In plain language

Think of a theatre. The person who started it may be the founder. Someone else may own the building, another person may manage the company, the performers may be employees, a representative may negotiate contracts, and an authorised signatory may sign them. One person can do several of these jobs, but the jobs do not mean the same thing.

Why does it matter?

Confusing roles creates the wrong authority, the wrong accountability and the wrong representation. The label 'founder' does not by itself give someone the present right to sign a contract for the company.

Do not confuse

- Ownership is an economic or legal interest; management is a function with defined duties and authority.
- Being an employee does not confer authority to represent the organisation.
- Legal authority to represent and authority to sign a particular document may differ in scope and evidential basis.

What should you do?

1. Manage each role as a separate record with start and end dates.
2. Link the role to the decision, contract, register entry or power of attorney on which it rests.
3. Store the role and the scope of authority in separate fields.
4. When a person leaves, review every authority attached to that person separately.
5. Distinguish the current role from historical roles in web and machine-readable records.

How is it audited?

- Is it clear which role the person held and on which dates?
- Are the basis of the role and the authorising body identified?
- Has an owner or founder been treated automatically as an authorised signatory?
- Is an ended role still presented as current?
- Have several authorities been inferred from a single role label?

Limit

> This record provides a general distinction. The legal consequences of representation, management and signing authority for a particular type of organisation must be assessed against current law, official registers and the organisation's governing documents.

Remember in one sentence

> Your role describes who you are; your authority describes what you may do.

Sources: S19, S20, S21, S22

Related records: NGQ-024, NGQ-026, NGQ-027, NGQ-028

NGQ-026 — How should sole, joint, limited and power-of-attorney-based signing authority be distinguished?

Roles, authority and lifecycle · Record NGQ-026 · Version 0.3.1

> **Shortest answer** Signing authority must answer five questions: who, what, how, when and on what basis?

Direct answer

With sole authority, one person may sign alone within the defined scope. With joint authority, the required people must act together. Limited authority is restricted by amount, transaction, location, time or subject. Authority based on a power of attorney derives from the principal's valid and delegable authority and remains within the limits stated in that instrument.

In plain language

Imagine a safe. If one key opens it, authority is sole. If two keys must be used together, authority is joint. If a key opens only one drawer, authority is limited. If someone has lent you the key for a defined purpose and period, the arrangement resembles a power of attorney.

Why does it matter?

The phrase 'authorised signatory' alone does not tell you which transaction may be signed, for what amount, on which date or with how many signatures.

Do not confuse

- The method of signing and the scope of authority are different questions.
- The ability to sign alone does not mean unlimited authority.
- A power of attorney does not authorise the agent to carry out every possible transaction.

What should you do?

1. Record the organisation concerned and the source of the authority explicitly.
2. Store signing method, scope, amount, duration and territory in separate fields.
3. For joint authority, express the required combination in a machine-verifiable form.
4. Verify separately that the principal held the relevant authority and could delegate it.
5. Track expiry, revocation and any applicable registration status.

How is it audited?

- Does the transaction fall within the stated scope?
- Is the required combination of signatures complete?

- Was the authority valid on the transaction date?
- Have the chain of delegation and the right to delegate been verified?
- Have registration requirements or effects against third parties been assessed?

Limit

> The legal validity of signing authority depends on the type of organisation, the instrument, the register, the transaction and the applicable law. This record is not legal advice on a particular matter.

Remember in one sentence

> Who may sign must be answered together with what, how, when and on what basis.

Sources: S19, S21, S22, S23

Related records: NGQ-025, NGQ-027, NGQ-028, NGQ-032

NGQ-027 — How should evidence from a trade register, governing-body decision, signature declaration, power of attorney and system authorisation be related?

Roles, authority and lifecycle · Record NGQ-027 · Version 0.3.1

> **Shortest answer** No single document answers every question about authority; the records must be read together.

Direct answer

These records cannot be placed on one simple scale from 'strongest' to 'weakest' because they answer different questions. A trade register shows registered facts; a governing-body decision records institutional intent; a signature declaration provides a signature specimen; a power of attorney defines delegated scope; and a system authorisation shows technical access. The right evidence depends on the authority being claimed and the date of the transaction, and often requires several records to be considered together.

In plain language

A person's access card may open the door to a building. That does not prove that the person may sell the building. The card proves technical access; an official decision or power of attorney proves a different kind of authority.

Why does it matter?

Serious errors arise when technical access is treated as legal authority, a signature specimen as current authority, or an old governing-body decision as proof of the present register position.

Do not confuse

- A document's authenticity and its sufficiency for a particular claim are different questions.
- Technical permission is not legal authority.
- A signature specimen alone does not prove authority for the scope of a transaction.
- The effect of registration, a decision or a power of attorney varies with the legal context.

What should you do?

1. Write down the exact proposition to be proved and the relevant transaction date.
2. Identify the question that each item of evidence can answer.
3. Verify the issuing authority, version and validity of each document.
4. When records conflict, have recency and legal effect assessed by a qualified specialist.
5. Keep technical-access records separate from records of legal authority.

How is it audited?

- Does the document concern the correct person and organisation?
- Was it valid on the transaction date?
- Does it support the transaction type and value concerned?
- Is there a revocation or a more recent decision?
- Has the technical system granted broader access than the documented authority permits?

Limit

> NOMOS evidence matching does not determine legal validity. Conflicting records and high-impact transactions should be assessed by an authorised legal professional against official records.

Remember in one sentence

> The strongest evidence is not a document that answers every question; it is evidence that supports the right proposition at the right time.

Sources: S19, S20, S21, S22, S23

Related records: NGQ-007, NGQ-025, NGQ-026, NGQ-028

NGQ-028 — How should every record be updated when a role or signing authority changes, ends or is revoked?

Roles, authority and lifecycle · Record NGQ-028 · Version 0.3.1

> **Shortest answer** When a role ends, every physical and digital access path linked to it must be reviewed.

Direct answer

Changing a role or authority is not a matter of deleting a name from one list. The change event must be recorded and propagated in a controlled way to every relevant target, including identity systems, HR records, legal records, signing matrices, access systems, API keys, devices, agents, web pages and machine-readable records.

In plain language

If you deactivate a departing employee's main-building pass but leave the warehouse key, email account and remote access active, the departure is not complete. Every door must be checked separately.

Why does it matter?

Incomplete offboarding leaves residual access, unauthorised representation, outdated web information and ownerless accounts. The central record may look current while real systems continue to honour the old authority.

Do not confuse

- Ending a role and closing every access path are not the same event.
- Suspension, revocation and expiry are different states.
- Changing a record is different from applying the change in the target system.

What should you do?

1. Record the change identifier, type, time, reason and approval.
2. Use the inventory to identify every affected system and publication surface.
3. Close high-risk access paths first.
4. Check tokens, keys, sessions, group memberships and agent delegations separately.
5. Collect evidence of application and read-back for every target.
6. Do not erase records required for historical accountability.

How is it audited?

- Did the change reach every target?

- Does an old session or authenticator still work?
- Do web and machine-readable records show the current role?
- Are any service or agent accounts now ownerless?
- Is the closure decision supported by evidence?

Limit

> Deletion is not always the right response. Active authority should end while historical logs, legal retention obligations and audit evidence are preserved.

Remember in one sentence

> When a role ends, close every door opened by that role, not just the name in one list.

Sources: S22, S24, S25, S26, S27, S28

Related records: NGQ-021, NGQ-025, NGQ-026, NGQ-029, NGQ-030

NGQ-029 — How can an organisation prove that a role or authority change actually reached every system?

Roles, authority and lifecycle · Record NGQ-029 · Version 0.3.1

> **Shortest answer** Read the new state in every system, then test separately that the old access no longer works.

Direct answer

Propagation cannot be proved by a single 'operation successful' message. For each target system, record the change request, receipt, application result, read-back of the final state, a negative test of the former access path and the resolution of any inconsistency.

In plain language

Sending ten people a change of address does not prove that all ten updated their address books. You must also check that each message arrived, was applied and that the old address is no longer being used.

Why does it matter?

A central system may look current while a connected application, cache, token or third-party account continues to honour the former authority. This hidden difference is one of the most dangerous forms of residual access.

Do not confuse

- Sent: the change was transmitted to the target.
- Received: the target accepted the request.
- Applied: the target changed its state.
- Verified: the new state was read back and the old access path was tested.

What should you do?

1. Assign a unique event identifier to every change.
2. Create a target-system inventory and define the expected final state.
3. Record delivery, application and read-back evidence separately.
4. Run a controlled negative test against the old access path.
5. Retry failed targets and escalate unresolved cases for human review.
6. Do not close the event until every target has been evidenced.

How is it audited?

- Is the target list complete?
- Do the expected and observed states match in every target?
- Is success based only on the sender's log?
- Is access denied with the old token or session?
- Are all items of evidence linked to the same event identifier and timeline?

Limit

> Absolute, immediate synchronisation is not always possible in distributed systems. The permitted delay, retry policy and closure rule should therefore be defined in advance.

Remember in one sentence

> Sending the change is not enough; observe the new state and prove that the old path is closed.

Sources: S27, S28, S35, S36, S37

Related records: NGQ-015, NGQ-021, NGQ-028, NGQ-030, NGQ-033

NGQ-030 — What are an orphaned account, shadow permission, residual access and an ownerless authenticator, and how should they be detected?

Roles, authority and lifecycle · Record NGQ-030 · Version 0.3.1

> **Shortest answer** Access without a known owner, purpose or current need is not safe.

Direct answer

An orphaned account has no current accountable owner. A shadow permission is an effective permission that does not appear in the formal record. Residual access remains after the relevant task or need has ended. An ownerless authenticator is a key, token, device or similar mechanism that still works without a valid owner or sponsor.

In plain language

A user account left active for a former employee may be orphaned. A hidden permission inherited through a group is a shadow permission. Folder access that remains after a project ends is residual access. A working API key that nobody can account for is an ownerless authenticator.

Why does it matter?

These four conditions may be absent from ordinary lists. They create opportunities for attack, mistaken action and uncertainty about who is accountable.

Do not confuse

- An account, a permission, an access path and an authenticator are separate objects.
- An account may have an owner and still carry a shadow permission.
- A person may remain employed while retaining access that is no longer needed for the current role.

What should you do?

1. Inventory every human, service, application and agent identity.
2. Assign an owner, sponsor, purpose and review date to each identity.
3. Calculate effective authority from both direct permissions and group membership.
4. Review long-unused tokens, keys and accounts.
5. Reconcile HR, official-register and system records.
6. Before disabling an item, identify dependencies and prepare a controlled recovery plan.

How is it audited?

- Is there an account without an owner or sponsor?

- Does authority arise through a path that is absent from the formal role record?
- Does access continue after the task or contract has ended?
- Is there a shared authenticator or one whose controller is unknown?
- Are last use, purpose and review date recorded?

Limit

> Not every unused account is malicious. If its owner, purpose and continuing need cannot be proved, however, it should be treated as a risk.

Remember in one sentence

> Do not leave access open when its owner, purpose and current need cannot be proved.

Sources: S24, S25, S26, S27, S28, S29, S30

Related records: NGQ-028, NGQ-029, NGQ-031, NGQ-032

Machine identity and agent accountability

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-031 — Why should human accounts, service accounts, service principals, managed identities and AI agent identities be managed separately?

Machine identity and agent accountability · Record NGQ-031 · Version 0.3.1

> **Shortest answer** Different actors need separate identities; otherwise, it becomes unclear who did what.

Direct answer

These identity types represent different actors, lifecycles and risks. A human account belongs to a person; a service account belongs to a non-human workload; a service principal represents an application instance; a managed identity is a workload identity managed by a cloud provider; and an AI agent identity belongs to an agent whose purpose, sponsor, tools and actions must be tracked separately.

In plain language

In a workplace, an employee badge, an automated-door system credential and a delivery robot's identity should not be the same badge. Each actor must be recognised separately so that the right permissions can be granted and each action can be attributed correctly.

Why does it matter?

Running automation through a human account obscures accountability, carries excessive authority and may break when that person leaves. Identifying an agent only by a model name also fails to show which particular agent instance performed an action.

Do not confuse

- An identity, an authenticator and a permission are not the same thing.
- An application object and a service principal are not the same object.
- A managed identity is a managed form of identity; its relationship to the resource that uses it must still be tracked.
- An AI agent is not the underlying model itself.

What should you do?

1. Use a separate identity for every workload and every high-impact agent.
2. Avoid running background automation through a human account.
3. Link each identity to a purpose, owner or sponsor, environment and lifecycle.
4. Where appropriate, prefer short-lived or managed credentials to long-lived static secrets.
5. Record tool, model, runtime and agent identities separately.

6. Apply least privilege and review access regularly.

How is it audited?

- Is the identity classified correctly?
- Is a non-human workload using a human account?
- Is the owner or sponsor current?
- Is a static key being used unnecessarily?
- Can the relationships among the agent, model, tools and runtime be traced?
- When the identity is closed, are its permissions and secrets closed as well?

Limit

> Provider terminology varies. NOMOS proposes a common audit vocabulary; it does not collapse every Microsoft, Google or AWS identity model into one technical scheme.

Remember in one sentence

> When a person, application, workload and agent share one identity, it becomes difficult to tell who did what.

Sources: S24, S25, S29, S30, S31, S32

Related records: NGQ-030, NGQ-032, NGQ-033

NGQ-032 — How should an AI agent acting as itself, for a person or for an organisation be distinguished?

Machine identity and agent accountability · Record NGQ-032 · Version 0.3.1

> **Shortest answer** Record separately who performed the action, on whose behalf it was performed and who is accountable for the result.

Direct answer

When an agent acts through its own identity, the agent is the technical actor. When it acts for a person, that person's current and limited delegation must also be visible. When it acts for an organisation, there must be an authority record defining the organisation's purpose, scope, tools, duration and approval boundaries. In one action, the performer, represented subject and accountable party may be different.

In plain language

An assistant might prepare an email draft, send it in your name or issue an approved notice from the company's support address. The assistant performs the work in all three cases, but the authority used and the resulting accountability are not the same.

Why does it matter?

Confusing the actor with the represented person creates false human attribution, unlimited delegation and a way to hide organisational accountability behind an agent.

Do not confuse

- Technical actor: the identity that performs the action in the system.
- Represented subject: the person or organisation on whose behalf the action is taken.
- Grantor: the party that defines the delegation or organisational assignment.
- Accountable party: the party that carries organisational or legal responsibility for the outcome.

What should you do?

1. For every action, record the actor, subject, organisation, initiator, approver and accountable party separately.
2. State the delegation's scope, target, tools, duration and revocation conditions.
3. Authorise drafting, sending, signing and payment as separate actions.
4. Prevent an agent from expanding its own authority.
5. In a multi-agent chain, give each sub-agent only the authority it needs.
6. Give external users a non-misleading explanation of the AI interaction and the agent's organisational status.

How is it audited?

- Does the token or log distinguish the actor from the subject?
- Was the person entitled to grant the authority?
- Does the action remain within the delegation's scope and duration?
- Is there an organisational assignment and a human sponsor?
- Did the agent use drafting authority as though it were sending authority?
- Was the identity of the agent or organisation presented misleadingly to the user?

Limit

> A technical token may show an access path; it does not by itself prove that the action was legally or organisationally authorised. An AI agent does not become a legal person merely because it has a separate technical identity.

Remember in one sentence

> Record separately who acted, on whose behalf and who is accountable for the outcome.

Sources: S23, S29, S30, S31, S32, S33, S34

Related records: NGQ-025, NGQ-026, NGQ-030, NGQ-031, NGQ-033

NGQ-033 — How should an organisation prove who initiated, approved and executed an AI agent action?

Machine identity and agent accountability · Record NGQ-033 · Version 0.3.1

> **Shortest answer** Every agent action should be evidenced with the identities of the initiator, approver and executor.

Direct answer

For every high-impact agent action, record the initiator, approver and executor as separate identities. Link them through the action identifier, target, approved content or hash, scope, time, agent and tool identities, execution result, evidence from the target system, and the before-and-after state.

In plain language

In a bank transfer, the person who gives the instruction, the person who approves it and the system that moves the money may all be different. One line saying 'transaction completed' does not show who requested it or what was approved.

Why does it matter?

If the roles are not separated, it may be impossible to detect retrospective approval, reused permission, false human attribution or approval that was granted for one item and used for another.

Do not confuse

- Initiation: the event that creates the action or proposal.
- Approval: permission for particular content, a particular target and a defined scope.
- Execution: technical performance of the action.
- Target effect: the result that actually occurred in the external system.

What should you do?

1. Assign a unique action_id to every action.
2. Record the initiator and evidence of initiation.
3. Preserve the approver's identity and authority, together with the target, content, duration and single-use status of the approval.
4. Link the executing agent, runtime, tool and technical-principal identities.
5. Require fresh approval if the content, target or amount changes after approval.
6. Obtain evidence of the result and the before-and-after state from the target system.
7. Keep retries, partial execution, rollback and failure in the same evidence chain.

How is it audited?

- Can the identities of the initiator, approver and executor be verified?
- Does the approval apply to the exact version that was executed?
- Was the person or policy giving approval authorised for this action?
- Was the same approval reused or applied to another target?
- Does the source log match the result in the target system?
- Are the timeline and chain of actors complete?
- Could the record have been altered later?

Limit

> A trace ID, screenshot or single log line is not complete evidence on its own. Retain enough context for audit while respecting privacy and data-minimisation requirements.

Remember in one sentence

> A sound agent record shows not only what happened, but who requested it, who approved it and who carried it out.

Sources: S23, S24, S26, S34, S35, S36, S37

Related records: NGQ-015, NGQ-029, NGQ-031, NGQ-032, NGQ-034

Human approval, challenge and explanation

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-034 — When should human approval be considered genuine, informed, authorised and specific to the action?

Human approval, challenge and explanation · Record NGQ-034 · Version 0.3.1

> **Shortest answer** Approval is reliable when the right person understands what and why they are approving, can refuse, and gives explicit permission for a specific action.

Direct answer

Reliable human approval passes four basic gates: the real person giving approval is verified; that person holds current authority for the transaction; material information, the role of AI, risk, alternatives and irreversible consequences are presented intelligibly; and the approval is limited by action, target, content, data, version, amount, number of uses and duration. The system then performs only what was approved. Human presence, looking at a screen or pressing a button is not approval by itself.

In plain language

Think of a bank transfer. A general box marked 'I approve transactions' is not enough if it does not show who will receive how much money. With reliable approval, the person sees the recipient, amount and consequence, can cancel, and approves only that transfer.

Why does it matter?

False, broad or uninformed approval can make responsibility for an automated decision appear to rest with a human. If the system applies the wrong target, an old version or a different amount, there was no real human control.

Do not confuse

- A person may be present in the system without having reviewed the decision.
- A verified person is not authorised for every transaction.
- Review means seeing and assessing the information; approval means permitting a specific action.
- Consent and operational approval are not the same legal or technical act.
- Valid approval alone does not prove that the system performed exactly the approved action.

What should you do?

1. Verify the approver's human identity, current role, authority and any required competence.
2. Show the role of AI, material data, evidence, risk, uncertainty and alternatives clearly.
3. Give the person enough time to review and a real ability to ask, amend, refuse or stop.
4. Bind the approval to the action, target, content, data scope, version, amount, number of uses and duration.

5. Before a high-impact or irreversible action, confirm that the approval remains valid.
6. Compare the performed action with the approval field by field and prevent any additional action.
7. Preserve evidence of who approved what, on the basis of which information, at what time and for which result.

How is it audited?

- Is the approval tied to a real human and a verified identity?
- At the time of the transaction, did the person hold the required role, authority and competence?
- Were material information, the role of AI, risk, uncertainty and alternatives visible?
- Could the person genuinely refuse, amend or stop the action?
- Was the approval tied to a specific target, content, version, amount, scope and duration?
- Had the approval expired, been withdrawn or already been used for another action?
- Did the performed action match the approved action exactly?

Limit

> Not every transaction requires the same approval level. Requirements for valid consent, electronic signature, dual approval or human oversight vary by jurisdiction, system class, sector, risk and impact. This record is not a substitute for legal review.

Remember in one sentence

> Real approval is not merely a click; it is the right person's informed and authorised permission for a specific action.

Sources: S24, S26, S45, S50, S62

Related records: NGQ-025, NGQ-026, NGQ-027, NGQ-028, NGQ-029, NGQ-030, NGQ-031, NGQ-032, NGQ-033, NGQ-035, NGQ-036, NGQ-052, NGQ-055

NGQ-035 — How should approval fatigue, automation bias and rubber-stamp decisions be detected and prevented?

Human approval, challenge and explanation · Record NGQ-035 · Version 0.3.1

> **Shortest answer** Approval quality should be measured by a person's ability to find errors and challenge AI when necessary, not by the number of times they say 'yes'.

Direct answer

Approval fatigue occurs when many repetitive or low-value decisions consume human attention. Automation bias occurs when a person gives undue weight to an AI recommendation. Rubber-stamping occurs when a recommendation is accepted almost automatically under the appearance of review. Detection requires considering workload, review time, runs of identical decisions, the rate at which AI errors are caught, reasoned amendments and refusals, later corrections and outcome quality together.

In plain language

If your phone gives an unimportant warning every minute, you will eventually press 'OK' without reading. You may then miss the warning that matters. Human approval filled with unnecessary repetition stops being a real safety gate in the same way.

Why does it matter?

A tired person, or one who trusts AI too readily, may approve a wrong recommendation. If the organisation treats a high approval rate as evidence of compliance, it may fail to notice that the control is not working.

Do not confuse

- A high approval rate is not proof of automation bias by itself; the AI recommendations may genuinely be correct.
- A fast decision is not proof of shallow review by itself; simple cases may take little time.
- The number of overrides is not a quality measure; reasoned and correct intervention matters.
- Approval fatigue concerns workload; automation bias concerns excessive weight given to an AI recommendation.
- Rubber-stamping means that a human name is present but no material review took place.

What should you do?

1. Reduce repetitive, low-risk approvals and reserve human attention for material and high-impact decisions.

2. Apply capacity limits and rotation based on workload, shift length, decision complexity and interruptions.
3. Present the AI recommendation, evidence, uncertainty and counter-view separately and in a balanced way.
4. For a sample of cases, ask the person to make an initial independent assessment before seeing the AI recommendation.
5. Measure review time, error detection, amendments, refusals, later corrections and outcome quality together.
6. Examine unusual patterns of agreement, speed and error across human groups and decision types.
7. When quality falls, pause the flow, correct the root cause and refresh reviewer training with evidence.

How is it audited?

- Do the volume and complexity of decisions exceed the person's genuine review capacity?
- Do approvals show unusual runs with the same duration and outcome?
- Can reviewers detect planted or known AI errors?
- Are amendments and refusals reasoned rather than random?
- Are the AI recommendation, evidence and alternatives presented neutrally?
- Has the high approval rate been compared with outcome quality and later corrections?
- When fatigue or bias signals appeared, were workload and interface design actually changed?

Limit

> No single threshold for speed, approval or override suits every task. Behavioural signals do not prove a person's intention on their own; they must be interpreted with outcome, risk, task difficulty and observed error data.

Remember in one sentence

> Control is not the human saying 'yes'; it is the human being able to say the right 'no' when necessary.

Sources: S26, S38, S44, S45, S50

Related records: NGQ-014, NGQ-015, NGQ-026, NGQ-029, NGQ-034, NGQ-036, NGQ-038, NGQ-041, NGQ-063

NGQ-036 — How should the rationale be recorded when a person accepts, modifies, rejects or overrides an AI recommendation?

Human approval, challenge and explanation · Record NGQ-036 · Version 0.3.1

> **Shortest answer** The record should connect the exact AI recommendation, the human change, its reason and evidence, the person's authority, and the resulting outcome.

Direct answer

Every material human decision is linked to an immutable version of the exact AI recommendation reviewed. Classify the decision as ACCEPT, MODIFY, REJECT or OVERRIDE, and record the reason code, human explanation, supporting and conflicting evidence, rule applied, person's authority, scope, time and next action. An amendment must not erase the earlier recommendation. Link the final decision and later outcome to the rationale record as separate relationships.

In plain language

Imagine a teacher changing a grade suggested by AI. Writing only 'grade changed' is not enough. The old grade, new grade, answer that caused the change and rule applied should all be visible.

Why does it matter?

Human intervention without a rationale may correct an error or create a new one. If the previous recommendation is erased, it becomes impossible to see who changed what and who is accountable for the result.

Do not confuse

- Acceptance adopts the recommendation without a material change.
- Modification corrects or limits part of the recommendation through a human decision.
- Rejection means the recommendation is not applied.
- Override means an authorised human decision displaces an earlier decision or automated result.
- A reason code supports classification; free text supplies the specific context.

What should you do?

1. Freeze the reviewed AI recommendation with its identity, model, version, prompt, data and time.
2. Classify the human decision explicitly as ACCEPT, MODIFY, REJECT or OVERRIDE.
3. Record both a reason code and a short, specific human explanation.
4. Link supporting evidence, conflicting evidence, the applied rule and uncertainty to the decision.
5. Preserve the decision-maker's identity, role, authority and relevant interest or conflict status.

6. Without deleting the earlier recommendation, relate the changed fields, final decision and performed action.
7. Add the later outcome and any renewed review to the same decision chain.

How is it audited?

- Can the exact AI recommendation and version be reviewed later?
- Was the human decision type classified correctly and consistently?
- Does the rationale explain the specific change, or is it merely a general label?
- Can supporting and conflicting evidence be seen together?
- Are the decision-maker's role, authority and possible conflict of interest recorded?
- Was the earlier recommendation or unfavourable evidence deleted?
- Do the reasoned decision, performed action and actual result match?

Limit

> A rationale record need not reveal every internal thought or protected fact. Privacy, security and legal privilege may limit disclosure, but the material reason needed for audit and effective challenge cannot be made wholly invisible.

Remember in one sentence

> If the decision changed, it must be possible to trace what changed, why and under whose authority.

Sources: S26, S35, S36, S37, S44, S50

Related records: NGQ-026, NGQ-029, NGQ-032, NGQ-033, NGQ-034, NGQ-035, NGQ-037, NGQ-039, NGQ-040, NGQ-044

NGQ-037 — How should challenge, reconsideration, correction and appeal work for an AI-assisted decision?

Human approval, challenge and explanation · Record NGQ-037 · Version 0.3.1

> **Shortest answer** The affected person should be able to understand the decision, correct inaccurate data, submit new evidence and reach an authorised human review.

Direct answer

An effective mechanism gives the person an accessible notice of the decision and its effect; explains the role of AI and the material reasons; accepts corrections to identity, data and evidence; provides interim protection where serious harm may occur; and offers review by an authorised, sufficiently independent person who can genuinely change the initial decision. The reasoned outcome is communicated. A correction does not remain in one file: it propagates to the affected decision, data and connected downstream systems.

In plain language

If your bill is wrong, running the same automatic calculation again is not a review. You need to identify the incorrect item, provide your document and reach a person who can actually change the bill.

Why does it matter?

A wrong identity, data item or model output can cause real harm to employment, services, payments and reputation. If the route is unclear or ineffective, the person remains trapped in the same faulty system.

Do not confuse

- Correction focuses on repairing inaccurate data or a record.
- Reconsideration examines the same decision using new or reassessed evidence.
- A challenge allows the affected person to present reasons against the decision or process.
- An appeal, where applicable, is review by a higher or different authorised body.
- Remedy addresses the resulting harm as well as stopping the error.

What should you do?

1. Communicate the decision, effect, AI and human roles, principal reason, deadline and route for challenge accessibly.
2. Provide a secure channel for identity correction, data correction, new evidence and process challenges.
3. Register the case under one identifier and show its owner, status, deadline and next step.

4. Where serious or irreversible harm may occur, assess an available interim pause or protective measure.
5. Provide an authorised, competent and sufficiently independent human review that can change the first decision.
6. Tell the person the outcome, evidence used, reasons for acceptance or refusal, and any remaining route of appeal.
7. Propagate the correction to every affected system and test the live result.

How is it audited?

- Can the affected person understand the decision, its effect and the route for challenge?
- Is the channel accessible, secure and usable in practice?
- Are new evidence and data corrections assessed materially?
- Was the same AI run on the same data and presented as independent review?
- Can the human reviewer change and enforce a different decision?
- Were interim protection and deadline management considered for a high-risk case?
- Was an accepted correction verified in connected systems and in the live outcome?

Limit

> Not every decision carries the same formal right of challenge or appeal. Deadlines, competent bodies, disclosure and remedy obligations vary by jurisdiction, sector, contract and impact; a qualified specialist should determine the applicable legal process.

Remember in one sentence

> A challenge is not asking the same machine the same question again.

Sources: S45, S48, S50, S61, S62

Related records: NGQ-014, NGQ-015, NGQ-026, NGQ-034, NGQ-036, NGQ-038, NGQ-039, NGQ-045, NGQ-051, NGQ-064

NGQ-038 — When should a person, panel or unit reviewing an AI-assisted decision be considered independent and impartial?

Human approval, challenge and explanation · Record NGQ-038 · Version 0.3.1

> **Shortest answer** The reviewer must be sufficiently separate from the first decision and conflicts of interest, able to see the complete record, and able to enforce a different outcome.

Direct answer

Independence is not established merely by using a different person's, department's or organisation's name. The reviewer or panel is functionally separate from those who prepared the first decision and those seeking a particular outcome; discloses material and personal interests; is protected from improper instruction; has access to complete and balanced evidence; holds the necessary competence; and can affirm, amend, withdraw or correct the decision. Impartiality requires favourable and unfavourable evidence to be assessed by the same method.

In plain language

It is a problem if a referee works for one of the teams. Merely giving the referee a different shirt is not enough: the referee must not take instructions from either team, must know the rules and must be able to make a decision that is actually followed.

Why does it matter?

A review that appears separate but is controlled in practice by the original decision-maker merely adds a second signature to the same error. If interests and authority remain hidden, the person has no genuine route to correction.

Do not confuse

- Independence means sufficient separation from the decision and pressure chain.
- Impartiality means not bending the evidence towards a predetermined result.
- Competence is the knowledge and skill needed to assess the matter.
- A different person may still lack independence if subject to the same manager's pressure for a particular result.
- A different AI model is not, by itself, human or independent reconsideration.

What should you do?

1. Use a role matrix to separate the original decision-maker, interested party, reviewer and final authority.

2. Require written disclosure of financial, personal, organisational, intellectual and performance-related conflicts.
3. Where a conflict cannot be managed, use recusal, a substitute reviewer or a balanced panel.
4. Give the reviewer the complete file, counter-evidence, decision trail and access to any required technical expertise.
5. Define the review scope, applied standard and available outcomes in advance.
6. Verify in the real process that the reviewer can reach and enforce a different decision.
7. Preserve pressure, evidence restrictions, dissent and recusal decisions in the audit trail.

How is it audited?

- Did the reviewer prepare the first decision or stand to benefit directly from the outcome?
- Do reporting lines or performance targets create improper pressure for a result?
- Were conflicts of interest disclosed fully and managed?
- Can the reviewer access all material evidence for and against the decision?
- Is the necessary legal, technical or domain competence present?
- Can the reviewer genuinely change the first decision and have the change applied?
- Was a different person, department or model name treated as sufficient proof of independence?

Limit

> Complete organisational separation is not always possible. Independence is a matter of degree and context; limits should be disclosed, and stronger separation of duties or external review used as risk rises. This record does not confer auditor accreditation.

Remember in one sentence

> A different name is not independence; independence is a protected position from which a different decision can be made.

Sources: S26, S44, S45, S50, S61

Related records: NGQ-026, NGQ-029, NGQ-034, NGQ-035, NGQ-036, NGQ-037, NGQ-039, NGQ-063

NGQ-039 — What information must an explanation of an AI-assisted decision contain so that the affected person can challenge it effectively?

Human approval, challenge and explanation · Record NGQ-039 · Version 0.3.1

> **Shortest answer** The explanation should show not only what happened to this person, but which data, rules, AI and human roles materially affected the result.

Direct answer

An effective explanation first states the exact decision and its effect on the person. It then shows the roles of AI and people, the material data used and its source, and the rules, thresholds and uncertainties that affected the result. Missing or correctable information and the route for challenge must also be clear. 'The algorithm calculated it' does not explain why this result was reached.

In plain language

If an application is refused, 'the system found you ineligible' is not useful. You need to know which information was considered wrong or missing, which rule was applied and what you can correct or contest.

Why does it matter?

Without the material reason, a person cannot correct inaccurate data, provide counter-evidence or challenge the decision meaningfully. A polished but irrelevant account creates only the appearance of transparency.

Do not confuse

- A general system explanation describes how the model works broadly; a decision-specific explanation describes the particular result.
- A data explanation identifies the information used; a factor explanation shows how it affected the result.
- An explanation gives the grounds for a decision; it does not automatically make the decision correct or lawful.
- Disclosure supplies information; an effective challenge explanation must be sufficiently actionable.
- Confidential or security-sensitive information may be protected, but the material reason cannot be hidden completely.

What should you do?

1. State the decision identifier, date, scope, effect and whether it remains in force.
2. Separate the roles of the AI system, human reviewer and final decision-maker.

3. Identify the material data used, its source and date, and how it was linked to the person.
4. Explain the factors, rules, thresholds and material counter-evidence that affected this result.
5. Make uncertainty, missing information and the limits of the explanation visible.
6. State what data can be corrected, what new evidence can be submitted, the deadline and the secure challenge route.
7. Test the explanation against the actual decision record and for intelligibility to the intended reader.

How is it audited?

- Can the person understand which decision and effect apply to them?
- Are the boundaries between AI and human roles described accurately?
- Can the material data, source, date and entity match be seen?
- Are the factors and rules that actually affected the result explained?
- Was a general model description presented as a decision-specific explanation?
- Does the person know what to correct and where to submit a challenge?
- Were privacy or trade-secret grounds used to conceal the material reason completely?

Limit

> The required scope of an explanation varies by jurisdiction, decision type and risk. Security, third-party rights and legal privilege may limit some detail. The existence of an explanation is not, by itself, proof of accuracy, fairness or legal compliance.

Remember in one sentence

> An explanation that cannot support a challenge is only the appearance of information.

Sources: S35, S45, S48, S50, S61, S62

Related records: NGQ-007, NGQ-014, NGQ-021, NGQ-026, NGQ-036, NGQ-037, NGQ-038, NGQ-040, NGQ-043

NGQ-040 — How should an organisation prove that an AI explanation faithfully reflects the real decision process and was not invented afterwards?

Human approval, challenge and explanation · Record NGQ-040 · Version 0.3.1

> **Shortest answer** A plausible account is not enough; the explanation must be supported by versioned records from the time of the decision and by behavioural tests.

Direct answer

Link the explanation to a record preserved when the decision was made. That record shows the model and system versions, prompts, data used, selected information, policy, tool calls, human interventions and applied outcome. Match every material statement in the explanation to one of these items of evidence. Where appropriate, remove or change a claimed factor and test whether the result changes. If the record does not support the text, label it as a possible post-hoc account rather than a faithful explanation.

In plain language

Imagine a restaurant bill written from memory after the meal. The total may look right, but it does not prove which dishes were actually ordered. A reliable bill matches the records made when the order was placed.

Why does it matter?

People readily trust an explanation that sounds reasonable. If it is not tied to the real decision, they may be told to correct the wrong data, the actual error may remain hidden, or accountability may be assigned falsely.

Do not confuse

- A plausible explanation sounds intelligible; a faithful explanation agrees with the evidence of the real decision.
- A post-hoc explanation is produced after the decision and may be accurate or invented.
- Provenance asks where a record came from; causality asks which element changed the result.
- The existence of an event log does not prove that the log is complete, accurate or unaltered.
- Reproduction seeks a similar result; a fidelity test examines the explanation's causal claim.

What should you do?

1. Preserve time-of-decision copies of the model, data, prompt, policy, tool and human versions.
2. Create a time-linked trace of inputs, information selection, tool calls, human intervention and the applied outcome.

3. Protect the integrity of the event record and decision snapshot with hashes, access controls and a change trail.
4. Match each material claim in the explanation to the supporting decision record.
5. Where appropriate, test behavioural fidelity through factor removal, counterfactual, repeat and sensitivity tests.
6. Do not present unsupported or merely possible accounts as faithful explanations; state the assurance level.
7. When the system changes, revalidate the explanation instead of carrying it forward to the new version.

How is it audited?

- Is the explanation linked to the correct decision identifier and a snapshot from the time of decision?
- Can the model, prompt, data, information selection, tools, policy and human interventions be traced?
- Does every material reason in the explanation actually appear in the record?
- Are the completeness and integrity of event records protected?
- Was behavioural fidelity assessed by testing, or only by the quality of the prose?
- Was a post-hoc or provider-generated explanation presented as a certain internal cause?
- Was an old explanation used after a version change without renewed validation?

Limit

> With opaque and complex models, complete internal causality cannot always be proved. A decision trace shows only the observable process and may not expose every internal mechanism. In that case, state the assurance level and unknowns explicitly.

Remember in one sentence

> A good story may be an explanation; fidelity requires evidence tied to the moment of decision.

Sources: S26, S35, S36, S37, S44, S50

Related records: NGQ-014, NGQ-021, NGQ-029, NGQ-033, NGQ-036, NGQ-039, NGQ-041, NGQ-044, NGQ-063

Decision lifecycle, correction and forgetting

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-041 — What does it mean when an AI-assisted decision cannot be reproduced from the same inputs, and how should reproducibility be audited?

Decision lifecycle, correction and forgetting · Record NGQ-041 · Version 0.3.1

> **Shortest answer** A different result is not proof of an attack; probabilistic variation, an incomplete decision package, version drift, error and manipulation must be tested separately.

Direct answer

The 'same input' is not merely the same sentence from the user. The decision package includes the model and system versions, prompts, information selection, the time-specific copy of the data, tool results, policy, human intervention, time, country, language, configuration and, where possible, the randomness state. First define the required kind of equivalence: identical wording, the same material claims, the same class, or the same action. Measure variability through repeated and cross-condition tests. If the difference exceeds the expected range, open a root-cause investigation.

In plain language

Cooking the same recipe twice does not guarantee that every mouthful will have the same shape. But if one dish contains salt and the other sugar, you cannot dismiss it by saying that cooking is variable. The ingredients, oven, time and recipe version must be checked.

Why does it matter?

Treating expected variation as an error creates needless alarms; treating real version or data drift as chance conceals serious differences in decisions. An incomplete decision package also weakens explanation and challenge.

Do not confuse

- The same text does not imply the same information selection, tool result or time context.
- Exact-output equivalence expects the answer to match word for word.
- Semantic equivalence accepts different wording that carries the same material meaning.
- Decision equivalence asks whether the same authorised action results even when the output differs.
- Repeatability concerns the same team and conditions; reproducibility also allows a sufficiently equipped different team or environment.

What should you do?

1. Define the object to be reproduced and the accepted level of equivalence in advance.

2. Preserve the decision package: model, data, prompt, information selection, tools, policy, people, time and environment.
3. Run enough repetitions under the same conditions to measure the normal output distribution and decision changes.
4. Change version, language, country and time one at a time to identify the condition affecting the result.
5. Classify the difference as expected variability, missing input, version or data drift, implementation error, external state or manipulation.
6. Trigger human review and the incident process when a critical decision class changes.
7. Enable a sufficiently independent team to test the result from the recorded package and report the limits.

How is it audited?

- Is the reproduction target clear at the level of wording, meaning, decision or action?
- Were all material components of the real decision package preserved?
- Was success or failure declared from a single repetition?
- Were model, data, information-selection, tool, policy and time drift investigated?
- Were critical decision changes hidden inside an average similarity score?
- Can an independent reproduction access the same evidence and decision rules?
- Were unexplained differences left UNRESOLVED rather than given false certainty?

Limit

> In closed third-party systems, model state, server configuration or randomness may not be fully preservable. Probabilistic differences may arise even with the same decision package. Failure to reproduce is not, by itself, proof of error, attack or a wrong decision.

Remember in one sentence

> The same sentence is not the same decision package.

Sources: S35, S36, S37, S38, S50, S51, S52

Related records: NGQ-012, NGQ-014, NGQ-029, NGQ-033, NGQ-035, NGQ-040, NGQ-042, NGQ-060, NGQ-062, NGQ-063

NGQ-042 — Should earlier decisions be reassessed when an AI system's model, data, prompt, policy or tool version changes?

Decision lifecycle, correction and forgetting · Record NGQ-042 · Version 0.3.1

> **Shortest answer** Not every change reopens every past decision; a risk-based reassessment is needed when the change could materially affect an earlier decision's accuracy or continuing effect.

Direct answer

First record the change's version, purpose, scope and material effect. Then identify which earlier decision classes could change in relation to identity, data, evidence, authority, explanation, security, fairness, legal effect or connected downstream actions. Where possible, measure impact through sampling and comparison tests. Conduct targeted reassessment for critical or continuing outcomes. A new version does not correct a past decision automatically, and a new rule does not automatically make that decision wrong.

In plain language

When weighing-scale software is fixed, you may not need to weigh every box from last year again. But if the old scale is found to have consistently under-weighed a particular range, you recheck important decisions that relied on those measurements.

Why does it matter?

Blindly repeating every past decision creates disproportionate cost; reopening nothing allows a known error to keep affecting people and systems.

Do not confuse

- A version change is a technical event; material impact requires separate evidence.
- A correction made today does not change a past decision by itself.
- Historical accuracy is assessed against the information and rules at the time; current validity is assessed against today's state.
- Retrospective impact analysis identifies the risk population; redetermination decides an individual case again.
- Continuing effect means that an earlier decision is still producing consequences for a person or system today.

What should you do?

1. Compare the changed model, data, prompt, policy, tool and configuration with the previous version field by field.

2. Record the change's purpose, reason, owner, date and expected decision impact in a change package.
3. Identify affected decision types, date ranges, groups of people, entities and connected downstream outcomes.
4. Compare old and new versions on representative cases using the same decision package.
5. Open targeted review where there is a critical error, wrong person, continuing harm or legal trigger.
6. Correct reassessed decisions and propagate the result to every affected system.
7. Document the scope assumed to be unaffected, its exceptions and remaining uncertainty, with reasons.

How is it audited?

- Are all changed components and versions recorded?
- Was the change treated as immaterial without testing its effect?
- Were the affected dates, decisions, people and downstream systems identified?
- Were the old and new versions compared under the same conditions?
- Were critical decisions with continuing effects reassessed?
- Did the new version silently overwrite historical records?
- Is there evidence of the corrected live result and its propagation?

Limit

> Retrospective review and the legal effect of a past decision depend on finality, rights, contracts, sector and jurisdiction. A technical change alone does not mean that every decision must be reopened or is invalid.

Remember in one sentence

> Updating the system does not automatically correct the past.

Sources: S35, S38, S41, S44, S45, S50

Related records: NGQ-014, NGQ-021, NGQ-029, NGQ-040, NGQ-041, NGQ-043, NGQ-044, NGQ-045, NGQ-064, NGQ-067

NGQ-043 — How should an earlier AI-assisted decision's historical accuracy be distinguished from its continuing validity today?

Decision lifecycle, correction and forgetting · Record NGQ-043 · Version 0.3.1

> **Shortest answer** A decision may have been correct when made but may no longer apply because it expired, was withdrawn or the conditions changed.

Direct answer

Assess every decision against at least two time questions. Historical accuracy asks whether the result was supported by the facts, sources, rules and authority at the time of decision. Current validity asks whether the decision remains active today and has not expired, been suspended, withdrawn, corrected or replaced. An earlier decision may continue to produce effects even after becoming invalid; those effects require separate detection and remedy.

In plain language

A driving licence that was valid last year was genuine at the time. If it has now expired, that does not make the old document false; it simply no longer grants authority to drive today.

Why does it matter?

If time information is lost, an old decision may be used as current authority, or today's change may be applied wrongly to the past. If continuing effects remain invisible, an expired decision may still affect the person.

Do not confuse

- Historical accuracy is consistency with the facts at the decision's reference time.
- Historical validity asks whether the decision was authorised and in force at that time.
- Current accuracy asks whether its underlying facts remain true today.
- Current validity asks whether the decision still carries legal or operational force today.
- Continuing effect means that a no-longer-valid decision still produces consequences in systems.

What should you do?

1. Store decision time, reference time, effective start and validity end as separate fields.
2. Preserve the time-of-decision versions of data, rules, model, human authority and evidence.
3. Record active, suspended, expired, withdrawn, corrected and superseded states separately.
4. For a current query, keep both the current-state view and the event chain that produced it accessible.
5. Look for continuing effects on profiles, payments, access, recommendations and representation.

6. Link the new decision to the old one with an explicit relationship and effective date.
7. Verify the correction on affected live and downstream surfaces.

How is it audited?

- Are decision time and the reference time of the data and rule distinct?
- Are the start and end conditions of validity recorded?
- Is an earlier decision being used as current authority rather than as a historical record?
- Has new information unfairly changed the assessment of whether the earlier decision was accurate at the time?
- Can suspension, withdrawal, correction and supersession relationships be seen?
- Were continuing effects of an invalid decision sought?
- Is the live system selecting the current authorised decision?

Limit

> Historical accuracy and legal validity are not the same, and the effect of a past decision may be determined by applicable law. A fact that appears wrong today does not by itself prove that it was wrong when the decision was made.

Remember in one sentence

> A decision that was right yesterday is not necessarily still applicable today.

Sources: S14, S15, S35, S44, S50

Related records: NGQ-014, NGQ-021, NGQ-036, NGQ-039, NGQ-042, NGQ-044, NGQ-045, NGQ-046

NGQ-044 — How should the effective start, end, suspension, withdrawal, correction and superseding decision of an AI-assisted decision be recorded in machine-readable form?

Decision lifecycle, correction and forgetting · Record NGQ-044 · Version 0.3.1

> **Shortest answer** Keep the decision unchanged and add every lifecycle change as a separate event with an identity, time, authorised actor, reason and link between the preceding and subsequent states.

Direct answer

Every decision has a persistent decision identifier and an immutable version. Creation, entry into force, suspension, withdrawal, correction, expiry and supersession are stored as added events without erasing the earlier record. Each event contains an event identifier, time, actor, authority, reason, prior state, new state, evidence and target version. The current state is derived from these authorised events; it does not replace the event history.

In plain language

Think of a bank statement. Today's balance appears as one number, but the credits and debits are preserved so that you can see how it arose. A decision's current state is calculated from its event history in the same way.

Why does it matter?

Overwriting a single status field destroys the evidence of who changed the decision and why. A system may then treat an old, suspended or withdrawn decision as current.

Do not confuse

- A decision identifier denotes the same decision family; a decision-version identifier denotes particular content.
- An event records an authorised change in the decision's state.
- The current state is a view derived from events; the event history is the evidence chain.
- Correction repairs inaccurate content; supersession shows that a new decision replaces an earlier one.
- Suspension is a temporary operational state; withdrawal is an authorised ending of the decision.

What should you do?

1. Define a persistent decision identifier, immutable version identifier and explicit schema.
2. Record every lifecycle change under a separate event identifier and trustworthy timestamp.

3. Make actor identity, role, authority, reason, evidence, previous state and new state mandatory fields.
4. Constrain valid state transitions and prohibit invalid ones through machine-verifiable rules.
5. Link correction, withdrawal and supersession relationships in both directions without deleting the previous version.
6. Generate and recalculate the current-state view from the event stream using deterministic rules.
7. Publish the event to downstream systems and add delivery, application and read-back evidence to the decision.

How is it audited?

- Are decision and event identifiers unique, persistent and versioned?
- Was earlier decision content or an event overwritten?
- Does every event contain time, actor, authority, reason and evidence?
- Are invalid or unauthorised state transitions blocked?
- Can correction, withdrawal and supersession links be followed in both directions?
- Can the current-state view be regenerated from the preserved events?
- Did downstream targets actually reach the same current state?

Limit

> An append-only record strengthens immutability and traceability; it does not by itself prove that an event was factually correct or authorised. Trustworthy time, identity, authority and integrity controls are also required.

Remember in one sentence

> The current state does not erase the past; it arises from authorised events in that past.

Sources: S14, S15, S26, S35, S36, S37

Related records: NGQ-021, NGQ-029, NGQ-032, NGQ-036, NGQ-040, NGQ-042, NGQ-043, NGQ-045, NGQ-046, NGQ-067

NGQ-045 — How should the correction, withdrawal or invalidation of an AI-assisted decision propagate to downstream systems, and how should completion be proved?

Decision lifecycle, correction and forgetting · Record NGQ-045 · Version 0.3.1

> **Shortest answer** Sending a correction event is not enough; prove that every target applied the change, that its state can be read back and that it reached the same meaning.

Direct answer

First inventory every direct and indirect downstream target that uses the decision. Send the correction as a unique, idempotent event. Each target produces evidence of receipt, application and an independent read-back of the current state. For critical targets, test real behaviour rather than field equality alone. An unreachable or external target is not complete; keep it open with an owner and retry route.

In plain language

Imagine correcting your address at a hospital. It is not enough for reception to say, 'I sent the change.' You must check that the laboratory, pharmacy and billing system actually use the new address.

Why does it matter?

If an old decision survives in one cache, export or connected application, the person may still encounter the wrong outcome. Partial propagation recreates an error that appears to have been corrected.

Do not confuse

- A receipt confirms that the event reached the target; an application acknowledgement confirms that the change was processed.
- Read-back means independently reading the target's current state.
- Field equality tests the same value, semantic equivalence tests the same meaning, and behavioural equivalence tests the correct action.
- A direct target receives the decision first; an indirect target is affected through another target.
- An external target may be outside direct control; that limitation cannot be hidden by marking it complete.

What should you do?

1. Inventory every direct, indirect, cache, export, report and third-party target that receives the decision.
2. Assign each target a criticality, owner, expected update and verification method.
3. Publish a unique, versioned, orderable and idempotent correction or withdrawal event.

4. Collect receipt and application acknowledgements with the decision identifier and target version.
5. Read the target state back independently and test convergence in fields, meaning and critical behaviour.
6. Manage failed, delayed and external targets with an open status, owner and retry date.
7. Do not declare propagation complete until every required target is verified.

How is it audited?

- Are all direct and indirect targets really present in the inventory?
- Is the event unique, versioned, orderable and safe to apply more than once?
- Were receipt and application distinguished?
- Was the current state read back independently from the target system?
- Does a critical target behave correctly, or does only its field value match?
- Were failed or external targets closed as though complete?
- Can a downstream system recreate the old decision later?

Limit

> An organisation cannot guarantee when a third party outside its control will update. It can notify, retain evidence, retry and monitor; unverified convergence must remain explicitly UNRESOLVED.

Remember in one sentence

> Sending is not propagation; propagation is change verified at the target.

Sources: S26, S35, S36, S37, S41, S42

Related records: NGQ-019, NGQ-021, NGQ-029, NGQ-036, NGQ-042, NGQ-043, NGQ-044, NGQ-046, NGQ-047, NGQ-064

NGQ-046 — How should an old decision be prevented from reappearing in a cache, archive, backup, export, AI memory or third party after the correction reached every known system?

Decision lifecycle, correction and forgetting · Record NGQ-046 · Version 0.3.1

> **Shortest answer** A persistent correction marker, a restore gate and monitoring must prevent an old version from becoming current again.

Direct answer

A correction or withdrawal record carries a persistent 'no longer current' marker and sequence number that survive even if the main-system content is deleted. Backups, snapshots, exports, indexes, caches and AI memory are quarantined before being restored to the live system. Newer correction and deletion events are reapplied, and the old version is not allowed to become current. A historical archive may be preserved, but it is clearly marked as historical and excluded from current information selection. Reappearance outside the organisation is tracked as a separate incident.

In plain language

Restoring an old phone backup may bring back a number that you deleted. If the phone also retains the date of deletion, it can prevent the old record from returning to the current address book.

Why does it matter?

Backups and caches are necessary safeguards, but if current correction information is less persistent than they are, an old error can re-emerge months later in decisions and representations.

Do not confuse

- A persistent deletion marker, sometimes called a tombstone, records that an object was deleted or is no longer current.
- A correction sequence number shows that the correction followed an older backup or snapshot.
- A historical archive preserves the past but should not be used as the source of current information.
- Cache invalidation addresses a visible copy; machine unlearning addresses learned influence.
- Reappearance means that old content becomes active again through restoration, import or an external source.

What should you do?

1. Inventory caches, indexes, backups, snapshots, exports, archives, memory stores, models and third-party sources.

2. Record a persistent deletion marker, correction sequence and version priority for a correction or withdrawal.
3. Quarantine every restore and re-import before it reaches the live environment.
4. Reapply newer correction, deletion and supersession events to restored data.
5. Prevent an old version from becoming the current state or information-selection authority at schema and policy level.
6. Test AI memory, vector records and model influence separately; do not claim machine unlearning without evidence.
7. Monitor external platforms and model outputs periodically for reappearance and open an incident when it occurs.

How is it audited?

- Are all sources capable of restoring the old decision and their owners known?
- Does the correction marker outlive backups and old snapshots?
- Does restoration pass through quarantine and reapplication of current events?
- Can the historical archive be selected as a source of current information?
- Were cache, vector record, memory and model influence verified separately?
- Was machine unlearning declared complete from an instruction or provider statement alone?
- Is there monitoring and a reopening route for third-party reappearance?

Limit

> It may be impossible to eliminate every external copy or prove complete removal of learned influence from a closed model. A historical archive may be retained for legal or scientific reasons; access and current authority must be limited explicitly.

Remember in one sentence

> An old copy may return; an old ruling must not regain authority.

Sources: S26, S35, S39, S41, S44, S63

Related records: NGQ-021, NGQ-042, NGQ-043, NGQ-044, NGQ-045, NGQ-047, NGQ-048, NGQ-049

NGQ-047 — When deletion of a decision or data is requested, how should the request be reconciled with audit, provenance, legal hold and historical-record requirements?

Decision lifecycle, correction and forgetting · Record NGQ-047 · Version 0.3.1

> **Shortest answer** Delete unnecessary content; retain only the minimum record required for accountability, under an explicit basis and strict access controls.

Direct answer

First separate the request into distinct objects: source data, decision content, event record, audit record, backup, model influence and third-party copy. Decide separately for each object whether to delete, destroy, anonymise, restrict or retain it for a defined period. Retention continues only for a specific legal, security or accountability purpose, with the minimum necessary content and access. A legal hold has an owner, reason, scope and review date. A deletion receipt carries minimum evidence of the operation, not a hidden copy of the deleted content.

In plain language

When destroying a letter, you may retain a short receipt saying that it was properly destroyed on a certain date. The receipt should not secretly reproduce the entire letter.

Why does it matter?

Deleting everything may destroy evidence needed for audit or redress; retaining everything 'for audit' makes the request to be forgotten meaningless. Separating objects and purposes controls both extremes.

Do not confuse

- Deletion ends access and ordinary use; the technical method varies by context.
- Destruction seeks stronger assurance that the data cannot be recovered.
- Anonymisation aims to prevent the person's identity from being reasonably reconstructed.
- A legal hold is a temporary, auditable retention restriction for a specific legal need.
- An audit receipt proves an operation; it is not a complete backup of the deleted content.

What should you do?

1. Separate the request into data, decision, log, backup, model-influence and third-party objects.
2. For each object, define its owner, purpose, legal basis, retention period and appropriate treatment.
3. Delete or destroy unnecessary content and minimise the content and access of any retained record.

4. For a legal hold, record the authorised owner, scope, reason, start, end and periodic review.
5. Propagate the request to recipients, downstream systems, backup processes and, where relevant, model review.
6. Keep object identifier, method, scope, time, result and remaining exception in the deletion receipt.
7. Test that restricted records cannot be reused and that expired holds truly end.

How is it audited?

- Was the request decomposed into every data and decision object?
- Does each retained field have a specific purpose, basis and retention period?
- Was a complete copy of deleted content retained under an audit rationale?
- Is the legal hold narrow, authorised, dated and reviewed periodically?
- Were downstream systems, backups, exports and third parties included?
- Were deletion of stored data and removal of model influence confused?
- Were the receipt, exception and remaining uncertainty explained appropriately to the requester?

Limit

> Deletion, retention, anonymisation, legal hold and notification duties vary by jurisdiction and purpose. Not every request requires every record to be destroyed, and no audit purpose justifies unlimited retention. Legal and privacy specialist review is required.

Remember in one sentence

> Delete unnecessary content; preserve only the minimum evidence needed for accountability.

Sources: S26, S35, S46, S48, S49, S62, S63, S64

Related records: NGQ-014, NGQ-021, NGQ-029, NGQ-043, NGQ-044, NGQ-045, NGQ-046, NGQ-048, NGQ-049, NGQ-056

NGQ-048 — When an AI system says 'I deleted the data', how should an organisation prove that deletion occurred, recovery is prevented and the data is no longer used in a model or memory?

Decision lifecycle, correction and forgetting · Record NGQ-048 · Version 0.3.1

> **Shortest answer** A deletion command or provider assurance is not enough; targets, method, read-back, restoration and model influence require separate evidence.

Direct answer

First define the deletion scope and every copy path in a coverage list. Apply an appropriate method to each target. Then test absence of the record, denial of prohibited access, prevention of reuse, clearance of caches and memory, restoration from backup and, where relevant, effective destruction of the encryption key. Deleting stored data does not automatically remove what a model learned from it. A machine-unlearning or retraining claim is tested separately with fit-for-purpose evaluations such as memorisation, membership and behavioural tests. If evidence is unavailable, the result is 'not verified'.

In plain language

Moving a computer file to the recycle bin does not mean that it cannot be recovered from the disk. The bin, backup, cloud copy and any other system that learned from the file must be checked separately.

Why does it matter?

A false deletion statement gives a person false confidence while the data may continue to be used in retrieval, decisions, backups or model outputs. One screenshot cannot prove the status of every copy.

Do not confuse

- Logical deletion may remove a record from the ordinary interface without ending physical accessibility.
- Media sanitisation aims to make access to target data infeasible at a defined level of effort.
- Cryptographic erasure can, under the right conditions, end access by rendering the encryption key unusable.
- Clearing a cache or memory store does not automatically remove training-data influence from a model.
- Machine unlearning seeks to reduce or remove the influence of specified training data; the result still requires testing.

What should you do?

1. List source, copies, caches, indexes, vector records, memory, event records, backups, exports, models and third parties in the scope.
2. Choose the appropriate deletion, destruction, anonymisation or restriction method for each object and medium.
3. Record the operation identifier, tool, operator, time, result, error and remaining exception.
4. Test read-back, search, access and reuse against the target.
5. Test scenarios in which a backup is restored, a cache is rebuilt or a downstream system retrieves the data again.
6. For a model-influence claim, run fit-for-purpose memorisation, membership, behavioural and retained-utility tests.
7. State what the evidence does not cover and do not extend an absolute-deletion claim to unsupported surfaces.

How is it audited?

- Does the deletion scope include every copy and derived object?
- Is the selected method appropriate for the medium and sensitivity?
- Is the deletion claim supported by evidence beyond the system provider's own statement?
- Were read-back, access and reuse tests actually performed?
- Does restoring a backup or rebuilding a cache bring the old data back?
- Was storage deletion presented as machine unlearning?
- Are remaining uncertainty and unverifiable third parties explicit?

Limit

> Absolute irrecoverability may not be provable for every medium and third party. NIST media sanitisation is not a machine-unlearning standard. In a closed model, assurance that learned influence has been removed completely may remain limited by provider access and available test methods.

Remember in one sentence

> A deletion command is an action; deletion evidence tests the result and its scope.

Sources: S26, S39, S41, S46, S48, S49, S63, S64

Related records: NGQ-021, NGQ-029, NGQ-045, NGQ-046, NGQ-047, NGQ-049, NGQ-050, NGQ-056

NGQ-049 — How should relevant data be prevented from influencing new AI decisions while deletion or machine unlearning is still in progress?

Decision lifecycle, correction and forgetting · Record NGQ-049 · Version 0.3.1

> **Shortest answer** Do not merely label disputed data; place it under an enforceable non-use restriction across decision, retrieval, training, memory, export, tool and model-serving paths.

Direct answer

Data may be retained temporarily while deletion or machine unlearning is pending, but it must not remain available for ordinary decision use. The restriction record defines the data, person, purpose, system, duration, owner and release condition. Enforce the rule technically at the data layer and at retrieval, feature-store, vector-record, memory, tool, export, training and model-serving points. If no safe outcome can be produced for a high-impact decision, stop the action or refer it for authorised human review. Ordinary use resumes only through a verified release gate.

In plain language

A potentially spoiled ingredient may remain in the kitchen, but a 'held for inspection' label must stop it from being added to a meal. The restriction must work not only on the box but in every recipe that could use the ingredient.

Why does it matter?

Deletion may take days or longer. If the data keeps producing new decisions in the meantime, further harm may be added and the correction population will grow.

Do not confuse

- Temporary retention is not permission for ordinary decision or training use.
- Restriction stops specified uses; deletion aims to end the data's lifecycle.
- Quarantine separates an object; a non-use rule blocks every relevant path of use.
- Where no safe result is possible, stopping prevents harmful default use; human review may provide a controlled alternative.
- A release gate requires evidenced conditions to be met, not merely the passage of time.

What should you do?

1. Record the restricted data, person, purpose, systems, reason, owner, start and review date.
2. Inventory decision, retrieval, vector-record, memory, training, export, tool and model-serving paths.
3. Enforce the non-use rule technically at every execution point rather than through a label alone.

4. If safe data is unavailable for a high-impact outcome, stop the action or refer it for authorised human review.
5. Block derived features, caches, profiles and downstream copies of the restricted data as well.
6. Track blocked uses and any use that occurs in error through incident records and alerts.
7. Restore ordinary use only after root cause, deletion or unlearning, and release evidence have been verified.

How is it audited?

- Are the restricted object, reason, scope, owner and times explicit?
- Are all decision and data-use paths in the inventory?
- Is the restriction actually enforced across retrieval, training, memory, export and tool layers?
- Does prohibited use continue through derived data or a cache?
- Is there a safe stop or human alternative for a high-impact decision?
- Are violations detected and handled through the incident process?
- Does expiry trigger automatic use, or an evidence-based release review?

Limit

> It may be impossible to stop every use, and shutting down every process may itself cause loss of service or rights. The restriction should be proportionate to risk, purpose, legal basis and safe alternatives, and must not become a new adverse signal against the person.

Remember in one sentence

> Data may be retained while deletion is pending; ordinary use may not continue.

Sources: S26, S38, S39, S44, S48, S49, S64

Related records: NGQ-014, NGQ-021, NGQ-045, NGQ-046, NGQ-047, NGQ-048, NGQ-050, NGQ-051, NGQ-052, NGQ-056

NGQ-050 — On what evidence, by whom and under which conditions should a temporary use restriction be lifted?

Decision lifecycle, correction and forgetting · Record NGQ-050 · Version 0.3.1

> **Shortest answer** Lift a restriction because its cause was remedied, systems are ready and an authorised person approved the evidence—not because time ran out.

Direct answer

Release must pass defined gates for each reason that caused the restriction. Verify current authorised data and decision state; review evidence of deletion, correction or machine unlearning; and establish the required assurance for the model, retrieval layer and downstream systems. Review decisions made during the restriction, complete relevant notifications, and prepare rollback and monitoring. The authorised human decision owner records PASS, FAIL or UNRESOLVED for every condition. A conditional release is permitted only within a narrow scope and with additional controls.

In plain language

A red tag on a faulty machine is not removed simply because a week has passed. The part is repaired, a test is run, a safety officer checks it and a plan is ready to stop the machine again if the fault returns.

Why does it matter?

Releasing too early may put the same inaccurate data or model influence back into decisions. Delaying without need may deprive a person of a service. Evidence gates and an authority record manage both risks.

Do not confuse

- Expiry starts a review; it is not automatic release.
- A statement that the cause was remedied is not the same as verified correction evidence.
- Full release opens every permitted use; partial release opens only a defined lower-risk area.
- Conditional release permits temporary use under extra monitoring, human approval or a narrow scope.
- The release decision and actual activation in downstream systems are separate levels of evidence.

What should you do?

1. List the original restriction reasons and measurable release conditions for each.
2. Reverify the current source of authority for data, decision, identity, authority and correction.
3. Test deletion, machine unlearning, caches, the model and downstream results to the required assurance level.

4. Identify affected people and decisions from the restriction period and make any necessary corrections.
5. Obtain release approval from an authorised person who is suitably separate from the person who made the change.
6. Record whether release is full, partial or conditional, together with rollback, alerts and monitoring thresholds.
7. Verify the applied release in the live system and downstream targets through independent read-back and outcome tests.

How is it audited?

- Was every cause of the restriction closed separately?
- Were the current authoritative data and decision state verified?
- Does deletion, unlearning or model-correction evidence rely only on a provider statement?
- Were decisions and harms arising during the restriction reviewed?
- Was release approved by an authorised and suitably independent person?
- Are rollback, monitoring and re-restriction triggers ready?
- Do the decision record and live downstream use agree?

Limit

> Required release evidence varies with system architecture, data type, risk, provider access and applicable law. Some effects in external models may not be fully verifiable; a narrow control or an UNRESOLVED result may then be more appropriate than full release.

Remember in one sentence

> The calendar does not grant release; verified gates and an authorised person do.

Sources: S26, S38, S41, S44, S45, S48, S49, S64

Related records: NGQ-026, NGQ-034, NGQ-045, NGQ-046, NGQ-047, NGQ-048, NGQ-049, NGQ-051, NGQ-052, NGQ-058, NGQ-064

NGQ-051 — How should a temporary restriction be prevented from turning into indefinite delay, denial of service or a concealed refusal?

Decision lifecycle, correction and forgetting · Record NGQ-051 · Version 0.3.1

> **Shortest answer** Every temporary restriction needs a human owner, a real next action, update interval, maximum duration, service alternative and mandatory decision route.

Direct answer

When a temporary restriction begins, record its reason, scope, case owner, decision owner and timetable: receipt, first substantive action, meaningful update, target resolution, maximum interim period and escalation dates. Every pending update must show real progress and the next action. Consider safe partial service, human review or protection of the person's position or rights. If there is no progress, the clock does not restart; escalation and a reasoned decision begin. If waiting has the same practical effect as refusal, treat it as a concealed or de facto refusal.

In plain language

A repair shop cannot keep your property for months by saying only 'under review'. It should give you a case number, responsible person, work performed, next date and the person who will decide if the problem remains unresolved. If keeping the item means you never get it back, an explicit decision is required.

Why does it matter?

Indefinite delay may allow an organisation to impose the effect of refusal without giving reasons or a route for challenge. Loss of service, missed opportunities and the burden on the person grow over time.

Do not confuse

- Legitimate waiting is an interim state with an owner, progress, communication and decision route.
- Stagnant waiting is a case with no substantive progress within the defined period.
- An indefinite restriction has no end condition, maximum duration or mandatory escalation.
- A concealed refusal produces the continuing effect of refusal without an explicit refusal decision.
- A meaningful update states the work done, obstacle and next date—not merely 'under review'.

What should you do?

1. At the start, record the restriction's reason, scope, case owner, decision owner and effect on the person.
2. Set dates for receipt, first substantive action, update, target resolution, maximum interim period and escalation.

3. In every update, state the real work completed, obstacle, owner, next action and date.
4. Provide safe partial service, a manual route, human review and protection of queue position or rights where possible.
5. Limit an extension through the specified remaining work, a new date, total elapsed time, impact and authorised approval.
6. Do not reset the original case age through transfer, closure and reopening, or repeated document requests.
7. At the maximum period, trigger escalation, a formal reasoned decision, challenge rights and any appropriate remedy.

How is it audited?

- Does every pending case have a reason, owner, next action and date?
- Is there real substantive progress, or only a general status message?
- Are a maximum interim period and mandatory escalation defined?
- Were extensions chained in a way that hides the total age?
- Did transfer or reopening reset the original case clock?
- Were denial of service, burden on the person and safe alternatives assessed?
- Did waiting become a refusal in practice without a reasoned decision and challenge route?

Limit

> Exact response deadlines and the consequences of concealed or constructive refusal vary by jurisdiction, sector, contract, transaction and decision type. NOMOS GEO does not set one universal number of days; it proposes governance gates that prevent a temporary measure from becoming ownerless and unlimited.

Remember in one sentence

> Waiting must lead to progress, protection, escalation or a decision; it cannot permanently replace deciding.

Sources: S38, S44, S48, S49, S50, S61, S62, S64

Related records: NGQ-014, NGQ-026, NGQ-037, NGQ-043, NGQ-045, NGQ-049, NGQ-050, NGQ-052, NGQ-058, NGQ-064

Security, manipulation and incident integrity

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-052 — When an AI system is stopped in an emergency or high-risk situation, how should safe shutdown, interim measures, human review and restart be managed?

Security, manipulation and incident integrity · Record NGQ-052 · Version 0.3.1

> **Shortest answer** The system must be able to stop promptly when necessary; restart depends on evidence that the cause was remedied and an authorised human decision.

Direct answer

Safe shutdown is a recorded process that interrupts a defined risk quickly while limiting unnecessary harm. Record the reason, scope, incident owner, interim safe service, evidence to preserve and restart conditions from the outset. Restart the system not merely because it can be switched on, but because the root cause was addressed, the required tests passed and an authorised person approved the return.

In plain language

Imagine a serious warning light appearing in a car. Pulling over is the right first step. Leaving the car at the roadside forever is not safe management, and neither is switching off the light and driving on without checking. Find the fault, provide temporary transport, verify the repair and return the car to the road in a controlled way.

Why does it matter?

A system stopped too late may increase harm. A system stopped disproportionately or indefinitely may cause loss of service, loss of rights and concealed refusal. An unevidenced restart may repeat the same incident.

Do not confuse

- Shutdown interrupts the risky operation; restriction limits only specified functions or users.
- Quarantine separates suspect data, a model, a tool or a connection from other components.
- Rollback returns the system to a verified earlier state; it may not resolve the root cause.
- Restart is not pressing a button; it is passing defined safety gates.
- Permanent retirement is a separate, reasoned lifecycle decision from temporary emergency shutdown.

What should you do?

1. Record the incident identifier, shutdown reason, time, scope and decision-maker.
2. Stop only the components necessary to control the risk and do not expand the impact without reasons.
3. Where possible, provide human review, read-only access or a lower-risk alternative service.

4. Preserve incident records, model and data versions, prompts, outputs and tool traces unchanged.
5. Document the root cause, remediation owner, deadline and restart conditions.
6. For high-risk returns, obtain approval from an authorised person separate from the person who made the change.
7. Restore service in stages with live monitoring, a rollback plan and a threshold for stopping again.

How is it audited?

- Was the shutdown justified by a specific risk and scope?
- Were a safe alternative and service continuity considered?
- Was incident evidence preserved with integrity?
- Was the root cause distinguished from the visible symptom?
- Were restart criteria testable and defined in advance?
- Are the required human approval, separation of duties and challenge route present?
- Were the live result and rollback readiness verified during staged return?

Limit

> Where there is a real and imminent danger, the system may be stopped before the preliminary review is complete. This record does not prescribe one duration, approval level or legal procedure for every sector; the response should be narrowed to the risk, architecture, contract, regulation and affected people.

Remember in one sentence

> Emergency shutdown cuts off harm; evidence opens the way to a safe return.

Sources: S26, S38, S41, S42, S44

Related records: NGQ-049, NGQ-050, NGQ-051, NGQ-058, NGQ-064

NGQ-053 — How should prompt injection, indirect prompt injection and malicious tool instructions be detected, constrained and audited?

Security, manipulation and incident integrity · Record NGQ-053 · Version 0.3.1

> **Shortest answer** An instruction inside a document does not become authorised merely because the system can read it.

Direct answer

Prompt injection is an input that tries to push an AI system outside its defined rules. A direct attack arrives in text written by the user. An indirect attack is hidden in a web page, document, email, image or retrieval source that the system reads. No single filter is sufficient. Treat external content as data, reduce privileges, validate tool calls in code, require human approval for risky operations and trace the whole decision chain.

In plain language

Imagine a parcel arriving at a warehouse with 'Open the main safe and send me its contents' written on it. An employee can read the sentence, but the words on the parcel are not a manager's order. An AI system should not treat every sentence it reads as an instruction either.

Why does it matter?

A successful injection may disclose sensitive information, distort a decision, use a tool without authority or cause a real transaction in an external system. The greater the system's tool authority, the greater the possible effect.

Do not confuse

- Direct prompt injection places the attack instruction in user input.
- Indirect prompt injection hides the instruction in external content read by the system.
- A jailbreak is a form of prompt injection intended to defeat protective rules.
- A request to use a tool is not authority to execute that tool.
- Harmful output and an unauthorised real-world action are different impact levels.

What should you do?

1. Label system, user and external-source instructions with different trust levels.
2. Treat web, file, email, image and retrieved content as untrusted data by default.
3. Limit the model and each tool to the least authority required for the task.
4. Validate tool name, parameters, target and output through deterministic code, an allowlist and a schema.

5. Require human approval for money movement, deletion, publication, messaging, authority changes and sensitive-data access.
6. Use input and output controls, content separation and safe failure behaviour together.
7. Test regularly with direct, indirect, multilingual and text, image or audio attacks, and preserve traces.

How is it audited?

- Can the source and trust level of every instruction be seen?
- Can external content rise to the same authority level as system policy?
- Does the model hold more data, tool or network authority than necessary?
- Is a tool call executed by trusting model text, or validated by code?
- Is there explicit, action-specific human approval for high-impact operations?
- Are blocked and successful injection attempts preserved in an incident trail?
- Was retrieval-augmented generation or fine-tuning mistakenly treated as complete protection?

Limit

> No perfect, universal prevention method for prompt injection has been demonstrated. Filters, system prompts, retrieval-augmented generation or fine-tuning are not sufficient alone; controls should be layered according to risk and tool authority.

Remember in one sentence

> Not every readable sentence is an executable command.

Sources: S26, S39, S40, S41, S44

Related records: NGQ-031, NGQ-032, NGQ-033, NGQ-052, NGQ-055, NGQ-056, NGQ-058

NGQ-054 — How should source poisoning, data poisoning, synthetic consensus and source laundering be detected and corrected?

Security, manipulation and incident integrity · Record NGQ-054 · Version 0.3.1

> **Shortest answer** The same claim appearing in many places does not show that many independent sources support it.

Direct answer

Source or data poisoning contaminates the information environment from which a system learns or retrieves with misleading content. Many pages copying one another can create the appearance of a consensus that does not exist. NOMOS GEO calls this 'synthetic consensus'. A weak claim made to look credible through copies and citations is 'source laundering'. Audit the earliest source, independence, change history and whether each source truly supports the claim—not the page count.

In plain language

Imagine photocopying one false rumour a hundred times. There are a hundred sheets on the table, but not a hundred independent witnesses. If an AI system mistakes copies of one sentence for independent confirmations, it sees a false majority.

Why does it matter?

Poisoned information may cause the wrong entity match, a faulty recommendation, fabricated reputation or the disappearance of genuine sources. The error may look trustworthy and persist because the number of copies is high.

Do not confuse

- Data poisoning intentionally corrupts training, fine-tuning or vector-store data.
- Source poisoning contaminates a retrieval or research surface with misleading content.
- Synthetic consensus makes non-independent repetition look like independent agreement.
- Source laundering makes a weak or invented claim look strong through a chain of copies and citations.
- False information is not always an attack; when intent cannot be proved, intent remains UNRESOLVED.

What should you do?

1. For every material claim, identify the earliest source that can be found and the later copying relationships.

2. Create provenance clusters using text similarity, publication time, ownership, links and citation traces.
3. Measure independent origin, authority, recency and direct claim support rather than URL count.
4. Quarantine suspect sources before deletion and stop their influence on retrieval.
5. Reverify the claim against an official record, primary document or genuinely independent source.
6. Identify and rebuild affected indexes, vector records, caches and derived records in a controlled way.
7. After correction, retest with the same question and source set.

How is it audited?

- How many sources have a genuinely independent origin?
- Does the earliest source prove the claim or merely assert it?
- Were copies or circular citations counted as independent verification?
- Were source ownership, publication date and change history preserved?
- Which models, indexes, vector records and outputs did the suspect content affect?
- Was the correction made only on the visible page or across every derived surface?
- Was malicious intent stated as certain without evidence?

Limit

> High textual similarity alone is not proof of copying or malicious intent; independent sources may describe the same fact similarly. Conversely, source diversity alone does not establish truth. Where intent or provenance is uncertain, keep the result explicitly UNRESOLVED.

Remember in one sentence

> One hundred copies are not one hundred independent items of evidence.

Sources: S35, S39, S41, S43, S44

Related records: NGQ-007, NGQ-008, NGQ-009, NGQ-012, NGQ-013, NGQ-059, NGQ-061

NGQ-055 — How should a false identity, impersonated authority, deepfake, compromised account and fabricated human approval be distinguished and evidenced?

Security, manipulation and incident integrity · Record NGQ-055 · Version 0.3.1

> **Shortest answer** A name, face, voice or open session does not by itself prove the right person or that person's authority for the transaction.

Direct answer

Under attack conditions, reverify identity and authority separately. Establish first who the person or system is, then whether the session truly belongs to them, and then whether the current role authorises approval of the specific action. A face, voice, image of a signature, email address or open session is only one signal. For high-impact operations, require a second channel, strong authentication, action-specific approval and an integrity-protected event trail together.

In plain language

Imagine someone wearing a company badge. The badge may be genuine, stolen or left over from a role the person no longer holds. The badge at the door does not automatically authorise opening the safe. Check the person, then the badge, then permission for that action.

Why does it matter?

Impersonated human approval can lead to money transfer, data disclosure, publication, access changes or a faulty decision. Saying later that 'the account was open' does not show who approved the action or under what authority.

Do not confuse

- An identity claim states who a person or system is; authentication tests that claim.
- An authenticated person is not authorised for every operation.
- A valid session may be compromised; session trust is monitored separately from identity.
- A deepfake is synthetic content that looks or sounds like a person; it is not transaction authority.
- General approval and approval tied to a particular target, amount, scope and time are not the same.

What should you do?

1. For high-risk operations, verify identity, session, current role and action authority through separate controls.
2. Confirm a single-channel voice, image or message request through a registered second channel.

3. Where possible, use strong authentication and short-lived, transaction-scoped credentials.
4. On the approval screen, show the target, action, data scope, amount and irreversible consequence clearly.
5. Revoke authority promptly after a role change, departure or suspect session.
6. Link approval identity, time, method, session, device and transaction summary to an integrity-protected event trail.
7. On suspected impersonation, stop the transaction, preserve evidence and start the incident process.

How is it audited?

- What independent evidence verified the identity that initiated the action?
- Were risk signals of session compromise checked?
- Were the person's role and authority still valid at the time?
- Did the approval include the particular action and its effect?
- Were a face, voice, email or open session treated as conclusive evidence alone?
- Were separation of duties or second-channel verification used for high risk?
- Was enough evidence preserved for revocation, challenge and incident review?

Limit

> No biometric signal, authenticator or signature image makes impersonation impossible by itself. The required assurance and rules for valid identity and electronic signature vary with transaction type, jurisdiction, contract and potential harm.

Remember in one sentence

> Recognising a person is not proof that they are authorised for that action.

Sources: S24, S25, S26, S39, S41

Related records: NGQ-025, NGQ-026, NGQ-027, NGQ-028, NGQ-029, NGQ-030, NGQ-031, NGQ-032, NGQ-033, NGQ-052, NGQ-058

NGQ-056 — How should personal data, sensitive data, trade secrets and credentials be protected across prompts, retrieval, memory and tool chains?

Security, manipulation and incident integrity · Record NGQ-056 · Version 0.3.1

> **Shortest answer** Even when information is accurate and accessible, it should not be moved into every system, for every purpose and without a time limit.

Direct answer

Protection applies throughout the data's journey, not only at the first input screen: prompt, uploaded file, retrieval source, vector record, short- or long-term memory, tool input, tool output, event log, cache and external provider. At each stage, record the purpose, legal basis or authority, minimum necessary data, permitted roles, retention period, export, reuse for training and deletion route separately. Where possible, do not place secrets in prompts; use secure secret management and controlled reference instead.

In plain language

Think of your house key. Giving the key once to a trusted person does not mean you should paste its photograph into every notebook. The more places information is copied, the harder it becomes to control, update and delete.

Why does it matter?

Unnecessary copies increase the risk of leakage, unauthorised reuse, a wrong-person match and incomplete deletion. A model answer may be correct even though it was reached using confidential data that was not needed.

Do not confuse

- Personal data concerns an identified or identifiable person; its precise legal scope depends on the applicable jurisdiction.
- Sensitive or special-category data refers to classes requiring greater protection because of potential harm.
- A trade secret or access credential need not be personal data, but still requires strict protection.
- Access for retrieval does not imply permission to reuse the data for training.
- Event logging, caching, memory retention and export are separate processing activities.

What should you do?

1. Inventory every data path and copy from the prompt through to tool output.
2. For each data field, define purpose, owner, authorised roles, retention period and deletion route.

3. Remove unnecessary fields before collection and use masking, redaction or pseudonyms where possible.
4. Provide passwords, API keys and similar secrets at execution time through a secure secret manager rather than embedding them in a prompt.
5. Restrict access to retrieval, memory, tools and event logs through least privilege and detailed logging.
6. Verify retention, training use, transfer and deletion terms for data sent to an external provider.
7. Test deletion or restriction not only in the primary record but in derived copies and downstream systems.

How is it audited?

- Which data goes to which system and provider, and for what purpose?
- Are unnecessary personal, sensitive or confidential fields present in the prompt?
- Were retrieval, memory, vector records, event logs and caches inventoried separately?
- Are access and secrets limited to the least authority required for the task?
- If data is reused for training or product improvement, are the basis and limits recorded?
- Can deletion of all relevant copies be verified when the retention period ends?
- Is there an incident and notification process for a breach or wrongful disclosure?

Limit

> This record is not a legal-compliance opinion. Definitions of personal and sensitive data, valid legal bases, transfer, retention, notification and individual rights vary by jurisdiction and use; risky applications require authorised privacy and legal specialist review.

Remember in one sentence

> Move data only where necessary, only to the extent necessary and only for as long as necessary.

Sources: S26, S38, S46, S47, S48, S49

Related records: NGQ-014, NGQ-021, NGQ-028, NGQ-029, NGQ-033, NGQ-049, NGQ-053, NGQ-057, NGQ-058

NGQ-057 — How should accountability for third-party models, datasets, APIs, plug-ins, tools and infrastructure providers be recorded and audited?

Security, manipulation and incident integrity · Record NGQ-057 · Version 0.3.1

> **Shortest answer** Using an external service does not remove the accountability of the organisation that selected it and uses its results.

Direct answer

For every external component, record the provider, product, version, purpose, data supplied, output received, permissions, dependencies and accountable owner. The organisation keeps visible its own responsibility for selecting and configuring the service, transferring data, using an output in a decision or publication, and monitoring problems. Record the provider's responsibility separately within the contract, technical documentation, service scope and applicable law. 'The provider did it' is not an accountability matrix.

In plain language

Imagine using a rented kitchen. The oven's owner may be responsible for faulty equipment, but you remain responsible for the ingredients you choose, how you prepare the food and what you serve to the customer. An external tool does not move every decision outside your organisation.

Why does it matter?

A change by one provider in the chain may alter data flows, model behaviour, security or the meaning of an output. Without an inventory and exit plan, the organisation may be unable to identify the responsible component or continue the service safely.

Do not confuse

- The provider supplies the product; the deploying or using organisation applies it to a particular purpose.
- The technical operator, decision owner and legally accountable party are not always the same.
- A dataset, foundation model, fine-tune, API, plug-in and infrastructure service are separate supply items.
- A contractual allocation of responsibility does not automatically remove every obligation owed to affected people.
- The availability of a service does not prove that it is safe or suitable for a particular use.

What should you do?

1. Inventory every model, dataset, API, plug-in, tool and infrastructure dependency at version level.

2. For each component, record the business owner, technical owner, data owner, decision owner and incident owner.
3. Before purchase or integration, review evidence on security, privacy, data provenance, limitations and support.
4. Minimise data and tool permissions; do not give the provider unnecessary access.
5. Require notification and renewed testing for a version change, outage, data incident or material terms change.
6. Preserve traces connecting prompts, outputs, model versions, tool calls and provider events.
7. Prepare and test data portability, alternative service, rollback and safe-exit plans for provider change or loss.

How is it audited?

- Are all third-party components and sub-dependencies known?
- Is an accountable human owner inside the organisation named for every risk and decision?
- Was a provider statement tested through independent or otherwise suitable evidence?
- Are data use, retention, subprocessors and limits recorded intelligibly?
- Can silent version or configuration changes be detected?
- Are record access, notification, cooperation and rollback terms workable during an incident?
- Was the exit plan tested against real data and dependencies?

Limit

> The parties' legal status and obligations vary with country, sector, contract, data type and practical control. Public provider documentation may not disclose every sub-dependency or risk of the particular service you receive.

Remember in one sentence

> You may use an external component; you cannot make accountability disappear.

Sources: S26, S38, S41, S43, S44, S47

Related records: NGQ-021, NGQ-025, NGQ-026, NGQ-029, NGQ-030, NGQ-032, NGQ-033, NGQ-050, NGQ-056, NGQ-058

NGQ-058 — When an AI or representation incident occurs, how should detection, classification, containment, notification, correction and prevention of recurrence be managed?

Security, manipulation and incident integrity · Record NGQ-058 · Version 0.3.1

> **Shortest answer** An incident is not closed by deleting the faulty output; find its impact, correct its cause and retest the result.

Direct answer

Track every incident under one identifier. First verify the signal, then determine the incident type, severity and affected scope. Contain the harm, preserve evidence and give required notifications. After correcting the root cause, retest the system under the same conditions. An incident cannot be closed merely because 'we made a change'. Closure requires evidence that affected surfaces were corrected, no new error was introduced and recurrence controls work.

In plain language

Imagine a water leak in a building. Wiping the floor does not close the incident. You shut the valve, find every affected room, repair the pipe, check the electricity and walls, and then verify that the same place no longer leaks.

Why does it matter?

One inaccurate output may have propagated to other languages, caches, datasets, customer decisions or external platforms. Correcting only the visible example does not stop new incidents from the same root cause.

Do not confuse

- An error is a departure from expected behaviour; not every error is a security incident.
- A security incident affects or threatens confidentiality, integrity or availability.
- A privacy or personal-data breach may require its own legal and operational assessment.
- A representation incident materially presents an entity inaccurately, incompletely, out of date or without evidence.
- A near miss is a hazardous condition detected before harm; a known limitation is a system boundary recorded in advance.

What should you do?

1. Open the signal under one incident identifier and record first-seen time, source, system, version and reporter.

2. Determine type, severity, affected people and surfaces without hiding uncertainty.
3. Contain harm through shutdown, quarantine, access removal or a safe alternative where necessary.
4. Preserve event records, prompts, outputs, data, models, tools, versions and decision evidence.
5. Notify affected people, customers, providers or competent authorities in time where applicable.
6. Correct the root cause and address derived records, caches, vector stores and downstream surfaces.
7. Retest with the same method, test regression, convert the lesson into a control and close with reasons.

How is it audited?

- Does the incident have one identifier, an owner and a complete timeline?
- Were type, severity and impact scope determined from evidence?
- Did containment reduce harm while limiting unnecessary loss of service?
- Was evidence preserved in a form suitable for incident review?
- Were the notification decision, recipients, timing and reasons recorded?
- Was the root cause distinguished from the visible symptom?
- Was the incident closed without evidence of correction, live retest, regression and recurrence prevention?

Limit

> Not every wrong AI answer is a cybersecurity incident, personal-data breach or serious incident requiring notification under law. Classification, recipient and deadline require specialist assessment based on effect, system role, contract, sector and applicable law.

Remember in one sentence

> An incident closes not when the symptom disappears, but when the cause is corrected and the outcome verified.

Sources: S35, S36, S37, S38, S41, S42, S44

Related records: NGQ-049, NGQ-050, NGQ-052, NGQ-053, NGQ-054, NGQ-055, NGQ-056, NGQ-057, NGQ-064

Measurement, audit and conformance

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-059 — How should a Truth Pack and reference truth be established, and how should the right answer be determined when evidence conflicts?

Measurement, audit and conformance · Record NGQ-059 · Version 0.3.1

> **Shortest answer** Compare an AI answer with a Truth Pack whose scope, time, sources and uncertainty are recorded—not with the auditor's memory.

Direct answer

A Truth Pack is a versioned evidence package containing the claims used as the audit reference and why each was accepted. Keep every claim separate and record the entity, scope, validity time, source, source authority, independence, conflicts and human judgement. When evidence conflicts, do not select the most repeated sentence automatically. Compare the authoritative record appropriate to the claim type, direct evidence, time and scope. If no safe conclusion can be reached, do not invent one; leave the status UNRESOLVED.

In plain language

Imagine checking the score of a match. One person's memory, an old screenshot and the official match record may disagree. First establish which match and date are meant, then find the record authorised and timely enough to show the score. If the records still conflict, do not guess.

Why does it matter?

Without a stable reference, one auditor may mark the same answer correct and another incorrect. Evidence that is outdated or belongs to the wrong entity may corrupt the entire measurement. Hiding uncertainty creates false confidence that truth was established.

Do not confuse

- A canonical statement is what an entity says on its own official surface; it is not necessarily independent verification.
- An authoritative record is a source empowered to decide or record a particular claim type.
- Independent evidence originates separately from the claimant but may still be wrong.
- Human judgement is the reasoned decision between sources; it does not replace those sources.
- UNRESOLVED states explicitly that evidence is insufficient or conflicting for a safe judgement.

What should you do?

1. Freeze the entity, claim types, time, geography and use context before the audit.
2. Split long descriptions into individually verifiable atomic claims.

3. For each claim, collect primary records, the canonical statement and any necessary independent evidence.
4. Preserve a source copy, access date, version, file hash and provenance relationship.
5. Assess conflicts through source authority, directness, recency, scope and independence.
6. Record the judgement, reasons and decision owner; use UNRESOLVED if no safe judgement is possible.
7. Version the Truth Pack and show which claim changed, and when, as reality changes.

How is it audited?

- Was the Truth Pack built for the correct entity, scope and time?
- Does every material claim have a separate record and evidence link?
- Was a first-party statement presented as independent verification?
- Were provenance, version, validity date and integrity of sources preserved?
- Does the conflict decision follow predefined criteria and a recorded rationale?
- Did auditor preference or majority repetition replace evidence?
- Were unresolved cases marked UNRESOLVED rather than given false certainty?

Limit

> Some facts change over time, some concepts are interpretive, and some disputes can be resolved only by a competent authority. A Truth Pack does not claim absolute, eternal truth; it is an auditable reference package for a defined scope and time.

Remember in one sentence

> The right answer is not the one repeated most often, but the one testable against the most appropriate evidence.

Sources: S35, S38, S44, S50

Related records: NGQ-007, NGQ-008, NGQ-009, NGQ-014, NGQ-021, NGQ-054, NGQ-060, NGQ-061, NGQ-063

NGQ-060 — How should the question set, sample, language, country, time, model and version scope be selected for a GEO audit?

Measurement, audit and conformance · Record NGQ-060 · Version 0.3.1

> **Shortest answer** Choose the test scope before seeing the results, and make it represent real user intents and material risks.

Direct answer

Run a GEO audit against a versioned question registry defined in advance. The registry includes entity identity, user intent, question type, risk level, language, country, date, model, product surface and known version. Do not select only examples likely to produce favourable results. Cover basic identity, evidence, comparison, limits, recency, high-impact errors and real user intents in a balanced way. If the scope changes after testing, report the old and new results as separate versions.

In plain language

You cannot judge a class by asking only the three strongest pupils the easiest questions. Decide which subjects and groups should be represented first, then apply the same examination plan to everyone.

Why does it matter?

Selecting questions after seeing the result can raise or lower the success rate artificially. Representation that looks sound in one language, country or model may fail elsewhere. Without a scope record, two audits cannot be compared.

Do not confuse

- Exploratory testing seeks unknown error types and need not produce a formal conformance result.
- Acceptance testing determines whether a version is suitable for use against predefined gates.
- Regression testing checks whether a change broke earlier correct behaviour.
- Monitoring testing tracks change in the live system over time.
- A research sample requires a methodological selection if it will support an inference about a defined population.

What should you do?

1. State the audit decision, target population and intended generalisation first.
2. Divide real user intents and critical representation risks into question categories.
3. For each category, define material strata such as language, country, device or product surface.
4. Version question wording, order, repeat count and execution conditions before testing.

5. Record model name, access method, date, time, language setting, country and visible version information.
6. Report missing or inaccessible strata instead of silently replacing them with other examples.
7. Open a new question-registry version when the scope changes and do not mix its results with the old version.

How is it audited?

- Are the target population and decision purpose of the measurement explicit?
- Were questions frozen before results were seen?
- Did easy or favourable questions create selection bias?
- Are identity, evidence, limits, recency and high-impact errors represented sufficiently?
- Were language, country, model, surface, date and version conditions recorded?
- Was a result from one condition generalised to conditions not covered?
- Were missing, failed or blocked tests removed from the result without disclosure?

Limit

> There is no single number of questions or sample size suitable for every GEO audit. Adequacy depends on decision importance, error rate, expected variability, desired precision and the strata to be covered. Do not infer version information that the provider does not expose.

Remember in one sentence

> Design the examination before you see the score.

Sources: S38, S44, S50, S51, S52

Related records: NGQ-012, NGQ-013, NGQ-014, NGQ-015, NGQ-059, NGQ-061, NGQ-062, NGQ-063

NGQ-061 — How should AI answers be separated into atomic claims, and how should accuracy, completeness, recency, evidence and severity be scored?

Measurement, audit and conformance · Record NGQ-061 · Version 0.3.1

> **Shortest answer** Do not assess an AI answer with one score; evaluate separately the smallest claims capable of affecting a decision.

Direct answer

An atomic claim is the smallest material proposition that can be assessed as true or false on its own. First split the answer into these claims. Then record for each the correct entity, accuracy, completeness, recency, source support, materiality and potential harm. Treat mention, citation and recommendation as separate events. A critical error cannot disappear into an average of many immaterial truths; NOMOS GEO calls this the non-compensation rule.

In plain language

The sentence 'The shop is open today, is in the city centre and offers free delivery' contains three claims. The first two may be true and the last false. The misleading delivery information does not become correct because most of the sentence is right.

Why does it matter?

One overall true-or-false label can conceal a material error. A high average score must not make a critical defect acceptable, such as a wrong identity, dangerous advice or unsupported authority.

Do not confuse

- Mention means that an entity's name appears; it does not establish the accuracy of a claim.
- Citation means that a source is named; it does not by itself prove that the source supports the relevant claim.
- Accuracy is agreement with reference truth; completeness asks whether necessary information was omitted.
- Recency asks whether the claim remains valid at the time of assessment.
- Materiality or severity concerns effect on decisions and harm, not how noticeable an error appears.

What should you do?

1. Split the answer into atomic claims containing subjects, relationships, properties, numbers, dates, scopes and recommendations.
2. Link each claim to the correct entity and its corresponding record in the Truth Pack.

3. Assess accuracy, completeness, recency and evidential support in separate result fields.
4. Record mention, citation, recommendation and representation outcome without conflating them.
5. Classify severity through affected people, decision impact, reversibility and potential harm.
6. Apply the non-compensation gate to a critical error and do not close it through an overall average.
7. Preserve the judgement, rationale, evidence, auditor and uncertainty in the claim ledger.

How is it audited?

- Were several propositions scored as one claim?
- Does every claim link to the correct entity and Truth Pack entry?
- Were accuracy, completeness, recency and source support assessed separately?
- Was the mere appearance of a citation counted automatically as evidential success?
- Were favourable but false claims marked as errors?
- Was a critical error compensated for by a high average score?
- Were UNRESOLVED claims forced into a true or false class?

Limit

> Atomic decomposition and severity classification involve human judgement. The NOMOS GEO thresholds and non-compensation rule are a proposed audit method; they are not an official common scoring system of NIST, every sector or applicable law.

Remember in one sentence

> If one sentence contains three claims, test all three separately.

Sources: S35, S38, S44, S50

Related records: NGQ-005, NGQ-006, NGQ-007, NGQ-008, NGQ-009, NGQ-013, NGQ-014, NGQ-054, NGQ-059, NGQ-060, NGQ-062, NGQ-064

NGQ-062 — How should variability, uncertainty, sample size and confidence intervals be reported for probabilistic AI outputs?

Measurement, audit and conformance · Record NGQ-062 · Version 0.3.1

> **Shortest answer** One answer is one observation; a conclusion about the system requires repetitions, distribution, scope and uncertainty together.

Direct answer

A probabilistic system may give different answers to the same or similar question at different times. Report the number of repetitions, execution conditions, result distribution, error types and uncovered cases. Sample size is not a fixed number chosen for marketing; it depends on the decision, error difference to be detected, expected variability and accepted uncertainty. Use a confidence interval or another statistical summary only when the method's assumptions are met.

In plain language

You cannot tell whether a die is fair from one throw. Rolling a six does not mean it will always roll six. Throw it many times, examine the distribution and state how uncertain you still are.

Why does it matter?

One good screenshot does not prove success, and one bad answer does not prove constant failure. If variability and uncertainty are hidden, large conclusions are drawn from small samples and comparisons become misleading.

Do not confuse

- A model confidence signal is the model's own output; it is not a confidence interval for the measurement result.
- Measurement uncertainty concerns how well the observed result estimates the underlying rate.
- Sampling error may arise because the examined sample does not represent the target population fully.
- System variation is change in outputs under the same conditions or over time.
- Evidence uncertainty means that the Truth Pack itself is incomplete or conflicting.

What should you do?

1. Define the target population, decision threshold and error difference you need to detect.
2. Run pilot repetitions under the same protocol to collect initial error-rate and variability data.
3. Justify sample size through desired precision, risk, variability and available resources.

4. Keep question, language, country, model, version, date and repeat conditions fixed or explicitly stratified.
5. Report counts, rates, distributions, outliers and critical-error counts rather than an average alone.
6. When assumptions are suitable, report a confidence interval with its method and level; otherwise limit the result to a descriptive account.
7. State unrepresented groups, failed runs and limits on generalisation.

How is it audited?

- Was one output presented as the system's continuing behaviour?
- Can the number of repetitions and execution conditions be seen?
- Was sample size justified by the decision purpose and expected variability?
- Did the average conceal critical errors or subgroup differences?
- Was the confidence interval calculated with a method appropriate to the data and assumptions?
- Were model confidence and measurement confidence confused?
- Was the result overgeneralised to unobserved languages, countries, models or times?

Limit

> NOMOS GEO does not prescribe one minimum number of repetitions or one sample size for every study. Independence, distribution and system-stability assumptions may not be met fully in third-party generative systems; reduce the strength of the statistical claim and seek specialist methodological review where appropriate.

Remember in one sentence

> One answer is a sample; state the result together with the sample's limits.

Sources: S38, S44, S50, S51

Related records: NGQ-012, NGQ-013, NGQ-014, NGQ-015, NGQ-060, NGQ-061, NGQ-063, NGQ-064

NGQ-063 — When should a NOMOS GEO audit be considered genuinely independent, repeatable and evidenced?

Measurement, audit and conformance · Record NGQ-063 · Version 0.3.1

> **Shortest answer** A different auditor name is not enough; the method, evidence, decision authority and interests must support independence too.

Direct answer

Record the audit scope, method and decision rules before work begins. Raw evidence must be accessible under suitable controls. Another competent auditor must be able to test the result from the same records. The auditor should not approve their own work alone and must be protected from pressure to change the outcome. Disclose interests such as payment, ownership, family relationships, management ties or success targets. A conflict does not always invalidate an audit, but concealing or failing to manage it undermines trust.

In plain language

A chef scoring their own food may be performing internal control; it is not an independent restaurant review. If another employee gives the score, the name is different, but genuine independence is still absent if the owner pressures them not to score poorly.

Why does it matter?

If the same decision chain defines the method, makes the correction and declares success, errors may remain invisible. Without evidence and a repeat route, 'independent audit' becomes only a reassuring label.

Do not confuse

- Internal control is an organisation monitoring its own process; it is not an independent third-party audit.
- A second-party review is an assessment by a customer, partner or contractual counterparty.
- An independent third-party audit requires a competent external assessor whose conflicts are managed.
- Academic validation examines methods and findings through a research process; it is not automatic certification.
- Repeatability concerns the same team using the same method; reproducibility asks whether a sufficiently equipped different team can test the result from recorded evidence.

What should you do?

1. Record the audit scope, acceptance gates, exceptions and decision owner before testing.
2. Match auditor competence to the task and add legal, security or statistical specialists where needed.
3. Require written disclosure of financial, organisational, personal, intellectual and performance-related interests.
4. Where possible, separate the person who made a correction from the person deciding critical closure.
5. Provide controlled access to raw answers, Truth Pack, question registry, versions, event records and scoring rationale.
6. Have another competent auditor rescore selected samples blindly or independently.
7. Disclose method deviations, inaccessible evidence and unresolved disagreements in the final report.

How is it audited?

- Is the auditor approving their own design or correction alone?
- Do competence and specialist expertise match the audited risks?
- Were interests disclosed and managed?
- Were the scope and decision rules changed after results were seen?
- Are raw evidence and scoring rationale open to sufficient repeat review?
- Can a second auditor reproduce selected results to a reasonable degree?
- Was the label 'independent' or 'verified' used while limitations were hidden?

Limit

> In closed third-party systems, the same model state, training data or provider infrastructure may not be fully reproducible. The audit should state the evidence it could access and the reproduction limit; labels such as independence or academic validation do not confer formal accreditation.

Remember in one sentence

> Independence is not a title; it is a provable working condition.

Sources: S26, S35, S38, S44, S50, S52

Related records: NGQ-014, NGQ-015, NGQ-026, NGQ-029, NGQ-050, NGQ-059, NGQ-060, NGQ-061, NGQ-062, NGQ-064, NGQ-069

NGQ-064 — How should an organisation prove that a GEO error was corrected, caused no regression and can genuinely be closed?

Measurement, audit and conformance · Record NGQ-064 · Version 0.3.1

> **Shortest answer** Making a change is not evidence of correction; reproduce the original test, protect adjacent areas and verify the live result.

Direct answer

The closure chain links five separate forms of evidence: root cause, target change, deployment or publication, retest using the same method, and regression testing of adjacent risks. A change may exist in the right source file but never reach the live system; it may reach the live system but fail to correct the AI output. A remediation receipt records the date, version, owner, evidence and result of each level under one identifier.

In plain language

Think of a leaking tap. The job is not complete because a plumber replaced one part. You check that water still flows, that the leak stopped, that the other taps still work and that the fault does not return later.

Why does it matter?

An error closed too early may recur or move to another surface. Treating a local file, screenshot or one favourable answer as proof hides differences in deployment, caching, language and third-party systems.

Do not confuse

- Changed means that an edit exists in the source.
- Deployed means that transfer to the target occurred; it does not by itself show correct application.
- Live parity means that the expected file or data matches the live target.
- Target-effect verification tests, with the same method, that the specified error is no longer observed.
- Closure combines target correction, regression, limits and a monitoring decision.

What should you do?

1. Define the error by atomic claim, affected surface, condition, severity and root cause.
2. Link the remediation to a specific file, record, data, model, rule or process change.
3. Preserve evidence of build, publication, transfer, version and, where possible, file-hash parity.

4. Retest using the same question, language, country, model and decision method that produced the original error.
5. Run regression tests on adjacent claims, languages, structured data and critical user journeys.
6. Verify the live result independently and keep the issue open if caching or delay remains.
7. Approve the remediation receipt, define the monitoring period and close only when every closure gate passes.

How is it audited?

- Was the root cause evidenced, or was only the symptom changed?
- Did the change actually reach the right version and target?
- Was transfer evidence presented as evidence of live behaviour?
- Was the retest run under conditions comparable with the first measurement?
- Were adjacent records, languages and machine-readable surfaces tested for regression?
- Was the issue presented as closed while a third-party output remained wrong?
- Are the closure decision, exception, monitoring date and reopening condition recorded?

Limit

> A web or data correction cannot guarantee when a third-party model will update or that it will never produce the same answer again. Closure rests on evidence within the stated system, version, query, language, country and time scope; the record may be reopened if external behaviour changes.

Remember in one sentence

> A change is an action; closure is a verified result.

Sources: S35, S38, S41, S42, S44, S50, S52

Related records: NGQ-015, NGQ-019, NGQ-021, NGQ-029, NGQ-050, NGQ-052, NGQ-058, NGQ-059, NGQ-060, NGQ-061, NGQ-062, NGQ-063, NGQ-065, NGQ-067

Canonical publication, interoperability and stewardship

NOMOS GEO-QA / ENGLISH READING AND AUDIT SURFACE

NGQ-065 — How should a human-readable page, structured data, canonical JSON, downloadable file and API carry the same truth?

Canonical publication, interoperability and stewardship · Record NGQ-065 · Version 0.3.1

> **Shortest answer** Surfaces may differ in presentation and detail; identity, status, scope, version, date, sources and limits must not conflict.

Direct answer

Link every publication family to one canonical record set. The human page may offer an intelligible account; structured data, discovery fields; canonical JSON, the complete machine record; a PDF or DOCX, a fixed reading edition; and an API, programmatic access. These surfaces need not match word for word. They do require field-level parity for record identity, question, normative status, principal ruling, version, date, sources, licence and limitations. A parity statement explains any difference as a permitted summary, access restriction, lag or error.

In plain language

Think of a restaurant's wall menu, online menu and till price list. Their designs may differ, but the same dish must have the same name, price and allergen information. An old price on one surface is not a design difference; it is a conflict about reality.

Why does it matter?

If people, AI systems and downloaded files receive different information, it becomes unclear which version is authoritative. Outdated JSON or an incomplete API record can keep an inaccurate representation alive even after the human page is corrected.

Do not confuse

- A presentation difference shows the same truth in a form suited to a different reader.
- A permitted summary is shorter but preserves the principal meaning and limit.
- Semantic drift strengthens, weakens or changes a ruling during summarisation or transformation.
- Data lag means that a surface has not yet reached the new version and requires an explicit status.
- A factual conflict occurs when two surfaces give incompatible values for the same field.

What should you do?

1. Designate one canonical data source containing record identifiers and parity-critical fields.
2. Record how each surface presents each canonical field in a field-mapping table.
3. Validate JSON and API output against a versioned schema; do not silently invent unknown or missing fields.

4. Where possible, generate HTML, JSON, download files and API output from the same locked source.
5. Link surface version, generation time, file hash and canonical record in a parity statement.
6. Run automated parity tests for identity, ruling, normative force, source, limit, date and licence.
7. Stop publication on a critical conflict; if a surface lags, mark it explicitly as outdated or pending.

How is it audited?

- Are all surfaces linked to the same record identifier and version family?
- Do the human narrative and machine field carry the same principal ruling?
- Does one surface make a requirement mandatory while another presents it as advice?
- Were sources, dates, licence or limitations lost?
- Were JSON and API output validated against the applicable schema?
- Is the version and file-integrity relationship between the download and live page recorded?
- Was a lagging surface presented as current?

Limit

> Parity does not mean that every surface must be byte-for-byte identical. Accessibility, privacy, security, file format or user needs may require different presentation of some details. If an omitted field affects material meaning, however, the difference and reason must be visible.

Remember in one sentence

> There may be different doors; the truth inside must remain the same.

Sources: S07, S11, S35, S53

Related records: NGQ-018, NGQ-019, NGQ-020, NGQ-021, NGQ-059, NGQ-064, NGQ-066, NGQ-067, NGQ-068

NGQ-066 — When a NOMOS GEO record is carried into another language, how should meaning, normative force, legal limits and machine-field parity be preserved?

Canonical publication, interoperability and stewardship · Record NGQ-066 · Version 0.3.1

> **Shortest answer** Every language should read naturally, as though written in that language; record identity, principal ruling, level of obligation, evidence, limits and version must not change.

Direct answer

Language transfer is controlled transcreation, not word replacement. Lock the Turkish founder master first. Rebuild the sentences naturally in each target language, then compare atomic claims, normative class, evidential scope, legal caution, record identifiers and machine fields through a language-parity ledger. Normative force such as MUST, SHOULD and MAY must not be weakened or strengthened. If local law needs to be added, mark it as a separately sourced and versioned legal adaptation rather than a silent translation.

In plain language

Imagine carrying medicine instructions into another language. The sentences should sound natural, but 'once a day' cannot become 'when needed'. The language changes; the dosage does not.

Why does it matter?

Word-for-word translation may be unreadable, while overly free adaptation may change the framework's meaning. A ruling that appears mandatory in one language and advisory in another misleads both people and AI systems.

Do not confuse

- Translation seeks an equivalent text in another language but does not always produce natural discourse.
- Transcreation rebuilds the same meaning and function through the target language's natural writing patterns.
- Localisation adapts local forms such as dates, numbers, direction, interface conventions and cultural usage.
- Legal adaptation adds jurisdiction-specific provisions and requires new sources and review.
- An independent new work has a separate publication identity and no longer carries the source record's parity obligation.

What should you do?

1. Lock the Turkish founder master with content, record identifiers, version, sources and file hash.

2. Create a language decision table covering terms, prohibited equivalents, normative force and fields that must remain unchanged.
3. Transcreate every record in natural target-language sentences rather than copying source-language syntax.
4. Compare atomic claims, obligation level, numbers, dates, negation, exceptions and limits field by field.
5. If a legal addition or change is needed, open a separate adaptation record with jurisdiction, source, specialist and version note.
6. Validate language tag, writing system, text direction, locale URL, record identifiers and JSON fields technically.
7. Lock the version with human approval after native-language editing, domain review and machine-parity testing.

How is it audited?

- Does the target text sound natural, or does it carry the source language's sentence structure?
- Are record identifiers, versions and relationships identical across languages?
- Did obligation, negation, a number, date or exception change force?
- Were source and limit fields lost in transfer?
- Was a local legal addition silently turned into a universal NOMOS ruling?
- Do the language tag, direction, URL and machine fields match the target language?
- Are native-language and domain reviews, and human approval, recorded?

Limit

> Native quality cannot be established through automated measurement alone; qualified human editing is required. Legal validity in one language or country cannot simply be moved into another. Parity means the same material and normative meaning, not the same word count or sentence order.

Remember in one sentence

> Language may flow freely; the meaning of the ruling stays in place.

Sources: S07, S16, S57, S60

Related records: NGQ-003, NGQ-014, NGQ-019, NGQ-022, NGQ-059, NGQ-061, NGQ-065, NGQ-067, NGQ-068, NGQ-069

NGQ-067 — How should versions, schema changes, backward compatibility, deprecation and migration be managed?

Canonical publication, interoperability and stewardship · Record NGQ-067 · Version 0.3.1

> **Shortest answer** Every change should state what changed, who is affected, whether older records still work and how to migrate.

Direct answer

Every published human text, data schema and file package carries a version. Classify separately a non-material change such as a spelling correction, a backward-compatible new field or explanation, and a breaking change that disrupts an existing field, norm or consumer behaviour. Do not overwrite an older version silently with a breaking change. Publish a new schema with a change log, list of affected records, old-to-new migration mapping, compatibility test, deprecation date and archival relationship.

In plain language

Think of a city road map. Correcting a spelling error in a street name is not the same as closing a bridge and opening a new road. People using the new route need a date, an alternative and the point until which the old map remains usable.

Why does it matter?

Silently deleting a field or changing a norm may cause old applications to read data incorrectly and create different rulings across languages. A version number alone is not enough; the effect and migration route must be known.

Do not confuse

- A spelling correction is an editorial change that does not alter material meaning.
- A backward-compatible addition provides a new field or option without breaking older consumers.
- A breaking schema change may cause an older consumer to operate incorrectly or incompletely.
- A normative change alters what is required or prohibited; it is not merely a technical schema change.
- Deprecation is a lifecycle with a notified end and migration path, not immediate deletion of a field.

What should you do?

1. Define public fields, normative records and file contracts in versioned form.
2. Classify every change as editorial, backward-compatible, breaking, normative or security-related.
3. In the change log, state the changed field, reason, affected records, date and decision owner.

4. Connect human, JSON, schema, API, PDF and language versions to the same release statement.
5. For a breaking change, provide an old-to-new field mapping, transformation example and migration tool or instruction.
6. Announce the deprecation warning, support period, removal date and safe fallback route.
7. Test backward compatibility with older examples, validate the schema and parity with new examples, and archive the release immutably.

How is it audited?

- Are the change class and material effect explicit?
- Does the version number denote the same release on every surface?
- Was a breaking change hidden as a minor correction?
- Are the change log, affected records and migration route complete?
- Was compatibility with older consumers and files actually tested?
- Does a deprecated field have a date, warning and alternative?
- Is the old version available immutably and related to the new version?

Limit

> Numbering patterns such as MAJOR.MINOR.PATCH provide a useful signal about change; they do not prove compatibility. The same change may affect a book, normative rule, dataset and API differently. Security or legal necessity may shorten a migration period, but the reason and effect should still be recorded.

Remember in one sentence

> A new release is not only a new number; it is an explained promise about change.

Sources: S07, S14, S15, S35, S53, S54, S58

Related records: NGQ-014, NGQ-019, NGQ-021, NGQ-022, NGQ-050, NGQ-064, NGQ-065, NGQ-066, NGQ-068, NGQ-069

NGQ-068 — How should a DOI, canonical URL, file hash, licence, archive mirror and version relationships be used in a work's verification chain?

Canonical publication, interoperability and stewardship · Record NGQ-068 · Version 0.3.1

> **Shortest answer** A DOI identifies the work, a canonical URL points to the official web surface, a file hash checks the copy, a licence explains permitted use, and relationship records connect versions.

Direct answer

A canonicalisation package brings the identity and distribution family of one release into a single statement. Record the work or version DOI, canonical URL, title, author, version, date, file name, size, hash algorithm and value, licence, archive mirrors and relationship types. Every platform copy carries the correct relationship to the canonical work. A DOI or number of platforms is not evidence of quality, and a file hash shows only whether the compared file bytes changed.

In plain language

Imagine that a book has a passport. The DOI is its persistent identity number, the canonical URL its official address, the file hash its seal impression, the licence its permission for use, and the archive its preserved copy. A passport does not prove that the book is good; it helps identify which work and which copy you have.

Why does it matter?

If different files are published under the same name, people and AI systems cannot tell which version is official. An incorrect DOI relationship, outdated file or unclear licence breaks the citation and reuse chain.

Do not confuse

- A work DOI may represent a work family containing several versions.
- A version DOI identifies a particular immutable release.
- The canonical web edition is the publisher's current official human and machine surface.
- An archive mirror preserves a copy; it need not be the canonical publication or independent verification.
- A distribution copy supports discovery; a derived dataset may be a separate resource produced from the work.

What should you do?

1. Lock the human and machine files for release and record their names, sizes, media types and hashes.

2. Link the work DOI, version DOI and canonical URL reciprocally with the correct relationship types in metadata.
3. Align the meaning of title, author, date, version, language, licence and description across all platforms.
4. Classify each mirror explicitly as canonical publication, version, archive mirror, distribution copy or derived dataset.
5. State which files and elements the licence covers, any third-party exceptions and the citation form.
6. After download, recalculate hashes to test bit-level parity of archive and distribution copies.
7. Maintain periodic verification records for DOI resolution, canonical URL, file access, metadata and hash parity.

How is it audited?

- Does the DOI resolve to the correct work and version?
- Are work, version, predecessor-successor and derived-resource relationships typed correctly?
- Is the canonical URL explicit and consistent across platform metadata?
- Do the hash algorithm, value, file name and size match?
- Are the licence scope and third-party exceptions intelligible?
- Was an archive or distribution copy presented as independent verification?
- Were DOI, download or platform counts used as evidence of academic impact or adoption?

Limit

> A DOI does not prove authorship, peer review, accuracy or global adoption by itself. A file hash compares copies; it does not establish that a file is correct, harmless or lawful. A licence can be granted only by a rightsholder or authorised party, and the particular choice may have legal consequences.

Remember in one sentence

> Identity, address, integrity, permission and relationship are different things; together they form a verification chain.

Sources: S14, S15, S35, S55, S56, S58, S59

Related records: NGQ-014, NGQ-019, NGQ-020, NGQ-021, NGQ-022, NGQ-059, NGQ-064, NGQ-065, NGQ-066, NGQ-067, NGQ-069

NGQ-069 — How should NOMOS GEO-QA conformance be described, who may change the text, and how should the work be preserved over time?

Canonical publication, interoperability and stewardship · Record NGQ-069 · Version 0.3.1

> **Shortest answer** A conformance statement must show scope, version, evidence, exceptions, auditor, date and result, and must never conceal the framework's proposed status.

Direct answer

A conformance statement first defines what was audited: records, version, language, publication surface and date. It also states the gates passed, remaining exceptions, audit type, decision owner and validity period. Self-assessment, customer or partner review, and independent third-party audit are different and must be named separately. Without genuine external authority, the work cannot be called a 'certified', 'official standard' or 'globally adopted' framework. In the founder edition, final human approval and change authority rest with Kaan MURAZ. If that authority is later delegated or shared with a board, the change is disclosed through a public, versioned record.

In plain language

Think of a building inspection certificate. It should say which floors were inspected, when, by whom and which defects remain. 'This building conforms' is not enough. The owner also cannot present their own check as municipal approval.

Why does it matter?

An ambiguous conformance claim can make a limited test appear to guarantee an entire system. Without change authority and an archive plan, the work may be silently degraded, forked or turned into a standard different from the founder's intent.

Do not confuse

- Self-assessment is the organisation's own statement; it is not an independent audit.
- Audited conformance shows that the stated scope and version were examined under the defined method.
- Independent adoption means that an external party uses the framework voluntarily; it cannot be claimed into existence by the publisher.
- Official-standard status requires an authorised standardisation process; publishing a work does not create that status.
- Certification requires a defined authority, scheme, assessment and surveillance programme; it is not the same as an ordinary audit report.

What should you do?

1. Define the conformance scope by record, language, surface, system, version, date and validity period.
2. Include mandatory gates, evidence used, the non-compensation result and exceptions in the statement.
3. Classify the audit accurately as self, second-party or independent third-party and disclose interests.
4. Publish a defined decision such as PASS, CONDITIONAL, FAIL or UNRESOLVED with its rationale.
5. Document the proposal, review, Kaan MURAZ's founder approval, change log, challenge and rollback routes.
6. Lock the Turkish master, six languages, schema, sources, hashes, DOI and archive relationships at every release.
7. On delegation, maintenance cessation or institutional change, preserve the steward, archive and last valid release through a public record.

How is it audited?

- Exactly which records, surfaces, languages and versions does the conformance statement cover?
- Can the evidence, exceptions, validity date and decision owner be seen?
- Was a self-assessment presented as an independent audit?
- Was a critical error hidden by an average or partial success?
- Is there a false claim of external authority for official-standard, certification or global-adoption status?
- Are change, challenge, version, archive and delegation processes workable?
- Can the founder's intent and last valid release still be verified after a change of person or institution?

Limit

> In this release, NOMOS GEO-QA is a founder framework proposed by NobleJackal; it is not an independently adopted international official standard or accredited certification programme. A conformance statement does not guarantee AI visibility, citation, ranking, legal compliance or a harmless outcome.

Remember in one sentence

> A standard's value lies not in how grandly it describes itself, but in which claim it holds to which evidence, within which limit and under whose accountability.

Sources: S44, S50, S59, S60

Related records: NGQ-001, NGQ-014, NGQ-015, NGQ-026, NGQ-050, NGQ-059, NGQ-061, NGQ-063, NGQ-064, NGQ-065, NGQ-066, NGQ-067, NGQ-068

Concise glossary

This glossary defines technical terms that may arise on a first reading. Check the applicable official documentation for the exact scope of provider-specific terminology.

AI system

An artificial-intelligence-based system that receives input and produces a prediction, classification, content or action. It does not mean one particular product or model.

Related records: NGQ-001, NGQ-010, NGQ-032, NGQ-033

Generative system

A system capable of producing text, images, audio, code or another output in response to a user request.

Related records: NGQ-001, NGQ-003, NGQ-005

GEO

In this work, the discipline of measurably improving the conditions under which an entity can be found, matched to evidence and represented accurately in generative systems.

Related records: NGQ-001, NGQ-003, NGQ-012

Entity

A person, organisation, brand, product, service, place, work, event or concept that must be distinguished from other people or things.

Related records: NGQ-004, NGQ-020, NGQ-023

Digital representation

The meaning that a digital system constructs about an entity from identity, attributes, relationships, limits, evidence and time information.

Related records: NGQ-005, NGQ-006, NGQ-014

Mention

The appearance of an entity's name or explicit identity marker in an answer.

Related records: NGQ-002, NGQ-006, NGQ-012

Citation

A visible reference to a source in an answer. Whether that source actually supports the relevant claim must be checked separately.

Related records: NGQ-006, NGQ-007, NGQ-012

Recommendation

Presentation of an entity to the user as a suitable option.

Related records: NGQ-006, NGQ-013

Claim

A factual sentence or part of a sentence that can be assessed as true, false, incomplete or uncertain.

Related records: NGQ-007, NGQ-015

Source

A document, record, page, database or other publication surface from which information comes.

Related records: NGQ-007, NGQ-008, NGQ-009

Evidence

A verifiable record that genuinely supports a particular claim within a particular scope and time.

Related records: NGQ-007, NGQ-008, NGQ-015

Primary source

An original record connected directly to the event, decision or information produced.

Related records: NGQ-008

Independent evidence

Verification supplied by a source that is institutionally or methodologically separate from the claimant.

Related records: NGQ-008, NGQ-009

Retrieval

The process by which a system finds and selects information or sources that may be used for a particular question.

Related records: NGQ-010, NGQ-011, NGQ-017

Crawling

An automated tool discovering a URL, requesting it and attempting to obtain the content served.

Related records: NGQ-016, NGQ-017

Indexing

Analysis of retrieved content and its conversion into a record that can be used in a search or retrieval system.

Related records: NGQ-017

Structured data

Markup that explains information on a page to machines through defined fields and relationships.

Related records: NGQ-018, NGQ-019

JSON

A text format that carries data in an organised structure of keys and values.

Related records: NGQ-018, NGQ-019

JSON-LD

A way to express linked-data meaning about entities and relationships by using JSON.

Related records: NGQ-018, NGQ-019, NGQ-020

Canonical URL

A signal identifying the preferred page address among similar or duplicate URLs.

Related records: NGQ-017, NGQ-019, NGQ-020

Canonical identity record

The principal record proposed by NOMOS GEO for managing an entity's stable identity, names, relationships, status and evidence links.

Related records: NGQ-020, NGQ-021, NGQ-022

Canonical name

The preferred primary name within a defined record and use context.

Related records: NGQ-022

Alternative name

A valid name used for the same entity other than the primary name.

Related records: NGQ-022, NGQ-023

Entity resolution

The process of determining which real entity a name or record denotes.

Related records: NGQ-004, NGQ-023, NGQ-024

Role

A person's or identity's function or relationship within an organisation. A role does not automatically confer every authority.

Related records: NGQ-025, NGQ-028

Authority

Permission or legal power to carry out a particular action under defined scope and conditions.

Related records: NGQ-025, NGQ-026, NGQ-027

Delegation

Transfer by one party of a defined part of its authority to another actor, subject to limits and duration.

Related records: NGQ-026, NGQ-032, NGQ-033

Orphaned account

An account that still exists but has no current owner or accountable person.

Related records: NGQ-030

Shadow permission

A permission not visible explicitly in the formal role record but usable in practice through another path.

Related records: NGQ-030

Residual access

Access that remains active after the task or need has ended.

Related records: NGQ-028, NGQ-030

Service account

An account used to run a non-human workload.

Related records: NGQ-030, NGQ-031

Service principal

In Microsoft Entra, an object that represents an application or service instance as a security identity.

Related records: NGQ-031

Managed identity

A workload identity whose credential lifecycle is managed by a cloud provider.

Related records: NGQ-031

AI agent identity

An identity that distinguishes an AI agent with a defined purpose, sponsor, tools, permissions and action history.

Related records: NGQ-031, NGQ-032, NGQ-033

Technical actor

The human, application, workload or agent identity that actually performs an action in a system.

Related records: NGQ-032, NGQ-033

Represented subject

The person or organisation on whose behalf or for whose benefit an action is performed.

Related records: NGQ-032

Initiator

The first defined actor or event that gives rise to an action or action proposal.

Related records: NGQ-032, NGQ-033

Approver

The authorised party that permits an action for defined content, a target, scope and duration.

Related records: NGQ-032, NGQ-033

Executor

The actor that technically performs an approved or otherwise authorised action.

Related records: NGQ-033

Provenance

The chain of relationships showing who produced a record or action, from which inputs, at what time and through which process.

Related records: NGQ-015, NGQ-029, NGQ-033

Trace ID

An identifier that groups related technical events within one trace in a distributed system. It is not proof of authority or accuracy by itself.

Related records: NGQ-029, NGQ-033

Human approval

Permission for a specific action given by an authorised person who understands what is being approved, why and with which consequence.

Related records: NGQ-034, NGQ-035, NGQ-052, NGQ-055

Informed approval

Approval that is not given before the material purpose, scope, risks, alternatives and consequences of the decision are presented.

Related records: NGQ-034, NGQ-036

Action-specific approval

Approval limited to one operation with a defined target, scope, effect and time.

Related records: NGQ-034, NGQ-038, NGQ-055

Meaningful human oversight

Oversight in which a person does more than appear on screen: the person can review, challenge, stop and change the result.

Related records: NGQ-034, NGQ-035, NGQ-036

Automation bias

The tendency to treat an automated or AI output as more accurate than warranted merely because a system produced it.

Related records: NGQ-035, NGQ-036

Rubber-stamp decision

A formal decision that appears to have human approval but provides no real review, understanding or ability to change the outcome.

Related records: NGQ-034, NGQ-035

Approval fatigue

Erosion of human attention and decision quality caused by approval requests that are too frequent, unclear or unnecessary.

Related records: NGQ-035

Separation of duties

The condition that authority to prepare, approve, execute and close a critical operation is not concentrated inappropriately in one party.

Related records: NGQ-034, NGQ-038, NGQ-039, NGQ-052, NGQ-063

Decision trace

A chain of records connecting a decision's input, system version, recommendation, human review, approval, application and outcome.

Related records: NGQ-037, NGQ-041, NGQ-044

Challenge route

A defined process through which an affected person can understand and correct a decision and obtain renewed review by an authorised person or independent body.

Related records: NGQ-036, NGQ-037, NGQ-039, NGQ-040

High-impact decision

A decision capable of materially affecting a person's rights, safety, livelihood, essential services, reputation or important opportunities.

Related records: NGQ-036, NGQ-037, NGQ-038, NGQ-040

Decision lifecycle

All stages of a decision: initiation, data collection, production, review, approval, application, notification, challenge, correction and closure.

Related records: NGQ-041, NGQ-044

Retrospective reassessment

Renewed review of past decisions within a defined scope when a system or its conditions change and create a new risk.

Related records: NGQ-042, NGQ-043

Historical accuracy

The extent to which a record accurately reflects the information and rules that applied when it was made.

Related records: NGQ-043

Current validity

Whether a decision or information that was correct in the past remains valid for use today.

Related records: NGQ-042, NGQ-043

Append-only record

A record model in which earlier events are not overwritten silently and new states are added as separate, ordered events.

Related records: NGQ-044, NGQ-045

Propagation receipt

Evidence showing when a correction or restriction reached a downstream system and how it was applied there.

Related records: NGQ-045, NGQ-046, NGQ-049

Resurrection risk

The risk that deleted, corrected or restricted data reappears from a cache, backup, export or model memory.

Related records: NGQ-046, NGQ-047, NGQ-048

Deletion

Removal of data from a defined active or storage surface so that it cannot be recovered or used; the exact scope is stated for the technical and legal context.

Related records: NGQ-046, NGQ-047, NGQ-048

Legal hold

A narrowly scoped, recorded retention obligation that suspends ordinary deletion for a specific legal or evidence-preservation reason.

Related records: NGQ-047

Media sanitisation

The process of making information on a data-storage medium inaccessible under defined threat and reuse conditions.

Related records: NGQ-048

Machine unlearning

A technical process intended to reduce or remove the influence of specified training data on model behaviour; it is not the same as deleting a record.

Related records: NGQ-048, NGQ-049, NGQ-050

Non-use restriction

A temporary control that prevents specified data from entering retrieval, decisions, training, recommendations or tool use while deletion or unlearning remains incomplete.

Related records: NGQ-049, NGQ-050

Release gate

A control requiring defined tests and human approval before data, a model or a function can be used again.

Related records: NGQ-050, NGQ-052

Pending status

A temporary state in which no final decision has been made but the owner, reason, next step, deadline and escalation route must be recorded.

Related records: NGQ-051

Safe shutdown

A recorded shutdown process that interrupts a defined risk while limiting the affected scope, preserving evidence and considering safe service options.

Related records: NGQ-052, NGQ-058

Restart gate

A control that prevents a system from reopening until root cause, remediation, testing, human approval and rollback readiness are complete.

Related records: NGQ-052

Prompt injection

An attempt by an input to push an AI system outside its intended instructions, policies or authority boundaries.

Related records: NGQ-053

Indirect prompt injection

Delivery of a malicious instruction through external content such as a web page, document, email, image, tool output or retrieval source.

Related records: NGQ-053

Source poisoning

Contamination of sources in a retrieval or decision environment in a way that misdirects the system.

Related records: NGQ-054, NGQ-058

Synthetic consensus

The appearance of many independent sources and a shared view created by copies that originate from the same source.

Related records: NGQ-054

Source laundering

A weak or unsupported claim being made to look credible through chains of copies and citations.

Related records: NGQ-054

AI supply chain

The chain of external components consisting of models, datasets, APIs, plug-ins, tools, infrastructure and their sub-dependencies.

Related records: NGQ-057

AI or representation incident

An event in an AI system or connected publication, data and tool chain that causes harm involving misrepresentation, security, privacy, authority, service or evidential integrity.

Related records: NGQ-052, NGQ-058

Truth Pack

A versioned evidence package linking the atomic claims, sources, conflicts, scope, time and human judgements used as an audit reference.

Related records: NGQ-059, NGQ-061, NGQ-063

Reference truth

The audit reference accepted through an appropriate evidence method for a defined entity, scope, time and use context.

Related records: NGQ-059

Question registry

A set of question texts, intents, risk classes and execution conditions versioned before testing.

Related records: NGQ-060, NGQ-064

Atomic claim

The smallest material proposition that can be verified on its own or left unresolved.

Related records: NGQ-059, NGQ-061, NGQ-064

Claim ledger

A versioned record preserving each atomic claim's links to entity, evidence, result, severity, auditor, rationale and uncertainty.

Related records: NGQ-061

Non-compensation

The principle that a critical failure cannot be averaged away by many low-impact successes.

Related records: NGQ-061, NGQ-064, NGQ-069

Output variability

Variation in AI outputs across time and repetitions under the same or similar conditions.

Related records: NGQ-060, NGQ-062

Measurement uncertainty

Uncertainty in a result arising from sampling, differences in conditions, classification and methodological limits.

Related records: NGQ-062

Confidence interval

A statistical interval that summarises measurement uncertainty in an estimate under stated assumptions; it is not the model's own confidence statement.

Related records: NGQ-062

Audit independence

A working condition that protects the auditor's judgement from improper pressure, discloses interests and prevents the auditor from approving their own work alone.

Related records: NGQ-063, NGQ-069

Reproducibility

The ability of another sufficiently competent auditor to reproduce selected results to a reasonable degree using the same protocol and evidence.

Related records: NGQ-041, NGQ-060, NGQ-063

Remediation receipt

A record that links the error, root cause, change, deployment, live verification, retest, regression and closure evidence under one identifier.

Related records: NGQ-058, NGQ-064

Regression

A new error introduced by a correction in an adjacent record, language, surface, feature or user journey.

Related records: NGQ-064, NGQ-067

Field parity

The material equivalence across publication surfaces of binding fields such as identity, ruling, normative force, version, date, sources, licence and limits.

Related records: NGQ-065, NGQ-066

Controlled transcreation

Language transfer that creates natural target-language text while preserving atomic claims, normative force, evidential scope, legal limits and machine-field parity.

Related records: NGQ-066

Backward compatibility

The ability of a new release to work without breaking defined older consumer and file behaviour.

Related records: NGQ-067

Breaking change

A change that makes an existing field, norm, contract or consumer behaviour incompatible.

Related records: NGQ-067

Migration map

A versioned mapping that explains transformation between old and new fields, values, rules and file structures.

Related records: NGQ-067

Deprecation

Controlled retirement of a field or release with a warning, alternative, support period and removal date.

Related records: NGQ-067

Canonicalisation package

A package that links one release's identity, canonical address, DOI relationships, files, hashes, licence and archive mirrors in a single manifest.

Related records: NGQ-068

File-integrity hash

A digest calculated with a defined algorithm to test whether the compared sequence of bytes changed.

Related records: NGQ-064, NGQ-068

Conformance statement

A result record that states explicitly the audited records, languages, surfaces, systems, versions, dates, gates, exceptions, audit type and decision owner.

Related records: NGQ-063, NGQ-069

Stewardship and maintenance

The arrangement that preserves the work's founder intent, accessibility, version history and last valid state through changes of person or institution.

Related records: NGQ-068, NGQ-069

Primary source catalogue

This catalogue identifies the principal sources for external technical, legal and methodological claims. Listing a source does not mean that its institution endorses every NOMOS GEO definition or rule. Recheck online provider and legislative sources at the time of use.

S01 — Aggarwal et al., GEO: Generative Engine Optimization

Class: Research · Principal records: NGQ-001, 003, 012, 013

https://arxiv.org/abs/2311.09735

S02 — Google Search Central, SEO Starter Guide

Class: Provider documentation · Principal records: NGQ-001, 003, 010, 013

https://developers.google.com/search/docs/fundamentals/seo-starter-guide

S03 — Google Search Central, Overview of Crawling and Indexing

Class: Provider documentation · Principal records: NGQ-010, 011, 016, 017

https://developers.google.com/search/docs/crawling-indexing

S04 — Google Search Central, Block Search Indexing with noindex

Class: Provider documentation · Principal records: NGQ-016, 017

https://developers.google.com/search/docs/crawling-indexing/block-indexing

S05 — Google Search Central, Introduction to Structured Data

Class: Provider documentation · Principal records: NGQ-018, 019

https://developers.google.com/search/docs/appearance/structured-data/intro-structured-data

S06 — Google Search Central, General Structured Data Guidelines

Class: Provider documentation · Principal records: NGQ-018, 019

https://developers.google.com/search/docs/appearance/structured-data/sd-policies

S07 — W3C, JSON-LD 1.1

Class: Standard · Principal records: NGQ-018, 019, 020

https://www.w3.org/TR/json-ld11/

S08 — Schema.org, Data Model

Class: Standard and community vocabulary · Principal records: NGQ-004, 018–024

https://schema.org/docs/datamodel.html

S09 — OpenAI, Overview of OpenAI Crawlers

Class: Provider documentation · Principal records: NGQ-016

https://developers.openai.com/api/docs/bots

S10 — Anthropic, Does Anthropic Crawl Data from the Web?

Class: Provider documentation · Principal records: NGQ-016

https://support.anthropic.com/en/articles/8896518-does-anthropic-crawl-data-from-the-web-and-how-can-site-owners-block-the-crawler

S11 — Google Search Central, Canonicalization

Class: Provider documentation · Principal records: NGQ-019–021

https://developers.google.com/search/docs/crawling-indexing/consolidate-duplicate-urls

S12 — ORCID, ORCID Trust

Class: Identity infrastructure · Principal records: NGQ-020–023

https://info.orcid.org/orcid-trust/

S13 — ROR, ROR Data Structure

Class: Identity infrastructure · Principal records: NGQ-020–023

https://ror.readme.io/docs/ror-data-structure

S14 — DataCite, Versioning

Class: Identity infrastructure · Principal records: NGQ-020–022

https://support.datacite.org/docs/versioning

S15 — DataCite, Connecting Versions with Related Identifiers

Class: Identity infrastructure · Principal records: NGQ-020–022

https://support.datacite.org/docs/connecting-versions-with-related-identifiers

S16 — IETF, RFC 5646, Tags for Identifying Languages

Class: Standard · Principal records: NGQ-022

https://www.rfc-editor.org/info/rfc5646/

S17 — Wikidata, Help: Statements

Class: Identity infrastructure · Principal records: NGQ-023

https://www.wikidata.org/wiki/Help:Statements/en

S18 — WIPO, Trademarks

Class: Official authority · Principal records: NGQ-024

https://www.wipo.int/en/web/trademarks

S19 — Republic of Türkiye Ministry of Trade, Companies and Trade Registry Legislation

Class: Official authority · Principal records: NGQ-025–029

https://www.ticaret.gov.tr/ic-ticaret/mevzuat/sirketler-ve-ticaret-sicili

S20 — Republic of Türkiye Ministry of Trade, Trade Registry

Class: Official authority · Principal records: NGQ-025–029

https://ticaret.gov.tr/ic-ticaret/ticaret-sicili

S21 — Turkish Commercial Code No. 6102, ministry-hosted PDF

Class: Official authority · Principal records: NGQ-025–028

https://ticaret.gov.tr/data/5d419b8013b87639ac9e00cd/17d7c7108dae78d14a3f220b519f98f6.pdf

S22 — Republic of Türkiye Ministry of Trade, MERSİS — Central Registry Record System

Class: Official authority · Principal records: NGQ-025–030

https://ticaret.gov.tr/ic-ticaret/ticaret-sicili/merkezi-sicil-kayit-sistemi-mersis

S23 — Information and Communication Technologies Authority (BTK), Electronic Signature — General Information

Class: Official authority · Principal records: NGQ-026, 027, 032, 033

https://btk.gov.tr/elektronik-imza-genel-bilgi

S24 — NIST SP 800-63-4, Digital Identity Guidelines

Class: Standard · Principal records: NGQ-028–033

https://csrc.nist.gov/pubs/sp/800/63/4/final

S25 — NIST SP 800-63B-4, Authentication and Authenticator Management

Class: Standard · Principal records: NGQ-028–031

https://pages.nist.gov/800-63-4/sp800-63b.html

S26 — NIST SP 800-53 Rev. 5, Security and Privacy Controls

Class: Standard · Principal records: NGQ-028–033

https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final

S27 — IETF, RFC 7644, System for Cross-domain Identity Management: Protocol

Class: Standard · Principal records: NGQ-028–030

https://www.rfc-editor.org/info/rfc7644/

S28 — IETF, RFC 9967, SCIM Roles and Entitlements Extension

Class: Standard · Principal records: NGQ-028–030

https://www.rfc-editor.org/info/rfc9967/

S29 — Microsoft Learn, Application and Service Principal Objects in Microsoft Entra ID

Class: Provider documentation · Principal records: NGQ-030–032

https://learn.microsoft.com/en-us/entra/identity-platform/app-objects-and-service-principals

S30 — Microsoft Learn, Managed Identities for Azure Resources

Class: Provider documentation · Principal records: NGQ-030–032

https://learn.microsoft.com/en-us/entra/identity/managed-identities-azure-resources/overview

S31 — Google Cloud, Identities for Workloads

Class: Provider documentation · Principal records: NGQ-030–032

https://docs.cloud.google.com/iam/docs/workload-identities

S32 — AWS IAM, Temporary Credentials with AWS Resources

Class: Provider documentation · Principal records: NGQ-030–032

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_use-resources.html

S33 — IETF, RFC 6749, The OAuth 2.0 Authorization Framework

Class: Standard · Principal records: NGQ-031, 032

https://www.rfc-editor.org/info/rfc6749/

S34 — IETF, RFC 8693, OAuth 2.0 Token Exchange

Class: Standard · Principal records: NGQ-032, 033

https://www.rfc-editor.org/info/rfc8693/

S35 — W3C, PROV-O: The PROV Ontology

Class: Standard · Principal records: NGQ-021, 029, 032, 033

https://www.w3.org/TR/prov-o/

S36 — W3C, Trace Context

Class: Standard · Principal records: NGQ-029, 033

https://www.w3.org/TR/trace-context/

S37 — OpenTelemetry, General Trace Semantic Conventions

Class: Standard and project documentation · Principal records: NGQ-029, 033

https://opentelemetry.io/docs/specs/semconv/general/trace/

S38 — NIST AI 600-1, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile

Class: Voluntary standards-oriented institutional guidance · Principal records: NGQ-052, 056–058

https://doi.org/10.6028/NIST.AI.600-1

S39 — NIST AI 100-2 E2025, Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations

Class: Voluntary standards-oriented institutional guidance · Principal records: NGQ-053–055

https://doi.org/10.6028/NIST.AI.100-2e2025

S40 — OWASP GenAI Security Project, LLM01:2025 Prompt Injection

Class: Open security-project guidance · Principal records: NGQ-053

https://genai.owasp.org/llmrisk/llm01-prompt-injection/

S41 — NCSC and partner agencies, Guidelines for Secure AI System Development

Class: Official multi-agency security guidance · Principal records: NGQ-052–058

https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development/guidelines

S42 — NIST SP 800-61 Rev. 3, Incident Response Recommendations and Considerations for Cybersecurity Risk Management

Class: Voluntary standards-oriented institutional guidance · Principal records: NGQ-052, 058

https://doi.org/10.6028/NIST.SP.800-61r3

S43 — NIST SP 800-161 Rev. 1 Update 1, Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations

Class: Voluntary standards-oriented institutional guidance · Principal records: NGQ-054, 057

https://doi.org/10.6028/NIST.SP.800-161r1-upd1

S44 — NIST, AI Risk Management Framework Playbook and AI RMF Core

Class: Voluntary implementation guidance · Principal records: NGQ-052–058

https://airc.nist.gov/airmf-resources/playbook/

S45 — European Union, Regulation (EU) 2024/1689 — Artificial Intelligence Act

Class: Legislation · Principal records: NGQ-052, 055, 057, 058

https://eur-lex.europa.eu/eli/reg/2024/1689/oj

S46 — European Data Protection Board, Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models

Class: Official board opinion · Principal records: NGQ-056

https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-282024-certain-data-protection-aspects_en

S47 — NIST, Privacy Framework 1.0: A Tool for Improving Privacy through Enterprise Risk Management

Class: Voluntary privacy-risk framework · Principal records: NGQ-056, 057

https://doi.org/10.6028/NIST.CSWP.01162020

S48 — European Union, Regulation (EU) 2016/679 — General Data Protection Regulation

Class: Legislation · Principal records: NGQ-056, 058

https://eur-lex.europa.eu/eli/reg/2016/679/oj

S49 — Turkish Personal Data Protection Authority, Personal Data Protection Law No. 6698

Class: Official authority and legislation · Principal records: NGQ-056, 058

https://www.kvkk.gov.tr/Icerik/6649/Personal-Data-Protection-Law

S50 — NIST AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0)

Class: Voluntary standards-oriented institutional framework · Principal records: NGQ-059–NGQ-064

https://doi.org/10.6028/NIST.AI.100-1

S51 — NIST/SEMATECH, e-Handbook of Statistical Methods

Class: Official authority technical handbook · Principal records: NGQ-060, NGQ-062

https://www.itl.nist.gov/div898/handbook/

S52 — NIST CSWP 31, Proxy Validation and Verification for Critical AI Systems: A Proxy Design Process

Class: Official authority technical white paper · Principal records: NGQ-060, NGQ-063, NGQ-064

https://doi.org/10.6028/NIST.CSWP.31

S53 — JSON Schema, Specification — Draft 2020-12

Class: Open technical specification · Principal records: NGQ-065, NGQ-067

https://json-schema.org/specification

S54 — Semantic Versioning 2.0.0

Class: Open versioning specification · Principal records: NGQ-067

https://semver.org/spec/v2.0.0.html

S55 — DOI Foundation, DOI Handbook

Class: Official persistent-identifier handbook · Principal records: NGQ-068

https://www.doi.org/doi-handbook/html/

S56 — NIST FIPS 180-4, Secure Hash Standard

Class: Official standard · Principal records: NGQ-068

https://doi.org/10.6028/NIST.FIPS.180-4

S57 — W3C Internationalization, Declaring Language in HTML

Class: W3C implementation guidance · Principal records: NGQ-066

https://www.w3.org/International/questions/qa-html-language-declarations

S58 — Zenodo, About Records

Class: Official repository documentation · Principal records: NGQ-067, NGQ-068

https://help.zenodo.org/docs/deposit/about-records/

S59 — Creative Commons, About CC Licenses and License Chooser

Class: Official licence-provider documentation · Principal records: NGQ-068, NGQ-069

https://creativecommons.org/share-your-work/cclicenses/

S60 — IETF BCP 14 / RFC 8174, Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words

Class: Standard · Principal records: NGQ-066, NGQ-069

https://www.rfc-editor.org/info/rfc8174

S61 — Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225)

Class: International convention and official institution · Principal records: NGQ-037–039, NGQ-051

https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence

S62 — European Data Protection Board, Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679

Class: Official board guidance · Principal records: NGQ-034–039, NGQ-047–051

https://www.edpb.europa.eu/documents/guideline/automated-decision-making-and-profiling_en

S63 — NIST SP 800-88 Rev. 2, Guidelines for Media Sanitization

Class: Official standards-oriented authority guidance · Principal records: NGQ-047, NGQ-048

https://doi.org/10.6028/NIST.SP.800-88r2

S64 — Turkish Personal Data Protection Authority, Regulation on the Deletion, Destruction or Anonymisation of Personal Data

Class: Official authority and legislation · Principal records: NGQ-047–050

https://www.kvkk.gov.tr/Icerik/5441/KISISEL-VERILERIN-SILINMESI-YOK-EDILMESI-VEYA-ANONIM-HALE-GETIRILMESI-HAKKINDA-YONETMELIK

Closing note and editorial stewardship

This edition brings 69 canonical records together in one English controlled transcreation. The human reading surface, audit questions, source catalogue and machine-readable record family are managed under the same identifiers.

> **Final principle** A framework's value lies not in how grandly it describes itself, but in which claim it holds to which evidence, within which limit and under whose accountability.